University of Southampton Research Repository

UNIVERSITY OF SOUTHAMPTON

Faculty of Engineering
School of Physical Sciences

Sparsification Techniques for Reduced Order Models of Turbulent Flows

by

Riccardo Rubini
PhD in Fluid Mechanics

A thesis for the degree of
Doctor of Philosophy

June 2022

University of Southampton

Abstract

Faculty of Engineering
School of Physical Sciences

Doctor of Philosophy

Sparsification Techniques for Reduced Order Models of Turbulent Flows

by Riccardo Rubini

The complexity of scale interactions, arising from the increasing number of dynamically active flow structures, is a well-known problem for the numerical modelling of high Reynolds number flows. Without doubts, this complexity is the main obstacle to the development of computationally affordable and physically interpretable models of complex flows. This research focuses on the nonlinear energy interactions across modes in reduced order Galerkin models of turbulent flows demonstrating a novel approach to automatically identify relevant interactions. This work is motivated by the key observation that, in the dynamics of high Reynolds number flows, not all the interactions have the same contribution to the energy transfer between flow structures. With the proposed work, we aim to develop a set of techniques to systematically select the dominant interactions in Galerkin models of turbulent flows, therefore identifying dominant triadic interactions. In the present work, two different approaches have been developed. First, a regression-based approach where the relevant interactions are identified *a posteriori* according to their relative strength. Second, an *a priori* approach, where a new set of basis functions, encoding the sparsity features of the flow, is generated. The key aspect of the latter approach is that the reduced-order model obtained by Galerkin projection onto the subspace spanned by the basis has sparse matrix coefficients without the need for any *a posteriori* evaluation. Both approaches have been tested on a set of flow configurations of increasing complexity. Results show that both approaches can identify the subset of dominant interactions preserving their physics throughout the sparsification process. In addition, further analysis showed that the *a priori* sparsification method preserves better the physics of triadic interactions, resulting in a better long term time stability and, therefore, should be preferred. Looking into the future, to scale up the *a priori* methodology to a more complex configuration some aspects need to be further investigated such as the role of the initial guess on the uniqueness of the result and its properties.

Contents

List of Figures

List of Tables

Declaration of Authorship

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as:

Signed:.. Date:..................

Definitions and Abbreviations

CFD	Computational Fluid Dynamics
DNS	Direct Numerical Simulation
LES	Large Eddy Simulation
ODE	Ordinary Differential Equations
PDE	Partial Differential Equations
MOR	Model Order Reduction
ROM	Reduced Order Model
POD	Proper Orthogonal Decomposition
ML	Machine Learning
XAI	eXplainable Artificial Intelligence
RNL	Reduced Nonlinear Model
DFT	Discrete Fourier Transform
FFT	Fast Fourier Transform
SPOD	Spectral Proper Orthogonal Decomposition
mPOD	Multiscale Proper Orthogonal Decomposition

Chapter 1

Introduction

Many phenomena in science and engineering are characterised by complex nonlinear interactions between structures at different spatial and temporal scales. In the case of fluid dynamics, the interactions between triads of fluid structures, called triadic interactions, play a fundamental role and have been shown to directly influence the physics of a wide range of different phenomena such as, energy cascade, transition to turbulence and many more. More specifically, the multiscale structure of the flow emerges from the structure of the triadic interactions and has two major implications mainly related to the study of the physics involved and to the computational efficiency of the mathematical models used. First, the computational resources required to simulate the flow grow quickly with the Reynolds number making a complete simulation of real-life configurations practically unfeasible. Second, the physical complexity influences the mathematical complexity making a satisfactory understanding of interactions between coherent structures often elusive. As a result, it is challenging to produce efficient and physically interpretable computational models of turbulent flows.

A possible solution that has received significant attention in recent years is model order reduction (MOR). This set of techniques aims to maintain the computational cost low preserving, at the same time, a good degree of physical fidelity. It consists in generating a physically meaningful and low dimensional representation of the flow, called reduced order model (ROM), by projecting the Navier-Stokes equations onto the subspace spanned by a suitable modal decomposition. Different methods to generate the basis vector defining this subspace are possible according to different properties of the problem considered such as geometry of the domain and physics involved. A complete overview on modal decomposition techniques for fluid dynamics is presented in Taira et al. (2017). However, one of the most popular approaches consists in projecting the governing equations onto the subspace spanned by the modes generated by the Proper-Orthogonal-Decomposition (POD) (Lumley, 1970, 1979). Historically, POD has been widely used for the generations of reduced order models for its energetic optimality that allows reconstructing a large amount of turbulent kinetic energy with the smallest possible number of modes. In addition, from a more physical standpoint, some authors suggested that POD modes can be a good mathematical approximation of coherent structures. Although this correspondence is not rigorous, POD can be a good tool to study energetic interactions between coherent structures in arbitrarily complex flows. Alongside several interesting properties, POD based reduced order models have also some drawbacks. The most relevant is related to the global mathematical nature of the spatial basis functions. It generates models that are mathematically dense (Noack et al., 2016), meaning that the evaluation of the dynamics of a single structure requires knowledge about the evolution of all the other structures in the flow. This feature leads to two major implications: the computational cost of complex models become quickly intractable and, consequently, the physical interpretation of the model itself becomes more and more challenging. These two issues can be mitigated by performing an early truncation of the modal expansion. However, this approach, equivalent to performing a coarser approximation of the flow, leads to inaccurate models with temporal stability issues and it is not a suitable solution. Another possible solution could be using different modal decomposition to perform Galerkin projection. However, only little work can be found in the literature regarding reduced-order modelling obtained by Galerkin projection on different subspaces than the POD, since the major effort has been directed into the mitigation of the issues of the POD based reduced-order models.

1.1 Relevant interactions in turbulence

This work is motivated by the established knowledge of the physics of energetic interactions in flow systems. The classical description of developed turbulent flows (Lumley, 1979; Pope, 2001; Jiménez, 2018) shows that energy is transferred across the hierarchy of coherent structures via nonlinear triadic interactions. Implicit in this picture is the fact that not all interactions have the same importance, but energy transfer occurs according to preferential patterns. This suggests that the nonlinear interaction pattern among coherent structures is sparse. The evolution of structures at a certain length scale depends predominantly upon a subset of all other structures (Kraichnan, 1971; Ohkitani, 1990; Brasseur and Wei, 1994) and the influence of interactions with the complementary set of structures can be generally neglected with minor global effects. This suggests that it would be possible to generate a mathematical model where, for a given set of flow structures or modes, only the interactions that actively contribute to the dynamics are considered. Sparsity features in triadic interaction are well documented for two and three-dimensional homogeneous turbulence. In addition, for three-dimensional configurations, the interactions have also been observed to be local, *i.e.* each structure interacts only with structures of commensurable size. Although this result is rigorously true only for homogeneous flows, some authors showed that qualitatively similar results are obtained in flow evolving in more complex geometries. More specifically, these features have been observed between coherent structures defined in a POD-based Galerkin model of a transitional boundary layer over a flat plate (see Rempfer and Fasel (1994a)). A similar pattern to the one observed for homogeneous turbulence has been also observed for the separated flow over a backward-facing step, as shown by Couplet et al. (2003).

The important result is that, regardless of the local/non-local nature of the triadic interactions, a subset of dominant interactions that regulate the energy transfer of the flow always exists if the reduced-order model has high enough dimensionality. Therefore, since not all the interactions contribute equally to the energy transfer it should be possible to generate models containing only the relevant interactions that are either computationally affordable and at the same time preserve the original physics. In the light of these physical results, this work aims to develop an automatic procedure to discern relevant interactions inside large Galerkin models of turbulent flows. The desired outcome is to obtain models containing a large range of spatial scales (a large number of modes) but where only the relevant terms for the dynamics need to be evaluated to temporally advance the model. In addition, because of the high complexity of the systems taken into consideration, it would be desirable that the procedure would be completely automatic without requiring any *a priori* knowledge about the physics of the system considered or external intervention by the user.

1.2 Model sparsification

Classically, modal basis functions used for model order reduction are generally obtained from considerations about energy optimality or the desired frequency content in each mode (Rowley and Dawson, 2017). However, they do not exploit the physical features such as sparsity of the interactions, since no information regarding nonlinear mechanisms that produce such interactions is utilised in the decomposition. For configurations without any particular symmetry when modal decompositions such as POD are employed, densely-connected models, *i.e.* models where all the entries of the matrices coefficients are different from zero, are usually obtained. This feature implies that to evaluate the the time advancement of the model all the possible terms of the model itself need to be computed. Crucially, the term governing the non-linear interactions in the Navier-Stokes equations is represented by a third order tensor whose computational complexity grows as the third power of the number of modes considered. This introduces some limitations for the application of this methodology to complex flow configurations. Develop a framework that can reduce this physical and computational complexity is the key enabler to extend the applicability of reduced-order models to more complex and relevant problems.

With model sparsification we intend a mathematical procedure that automatically identifies, within the reduced-order models, the terms associated with the dominant energy transfer, generating models where only the interactions that effectively contributes to the dynamics of the energy transfers are retained. The desired outcome would be to avoid truncation of the basis functions and obtain models of high accuracy, *i.e.* containing a wide range of spatial scales, but keeping, at the same time, the computational cost low, as only the necessary terms in the model need to be evaluated.

Successful attempts to construct reduced sets of equations, that exploit the sparsity of the energy interactions, have been made in the past, often for canonical geometries where triadic interactions are conveniently examined in Fourier space and using a coarse-grained partitioning of the hierarchy of scales. More specifically, this approach has been successfully applied to canonical flows such as two-dimensional homogeneous decaying turbulence and plane Couette flow (see Thomas et al. (2015)) showing that the reduced models can reproduce the fundamental statistical aspects of the flow.

However, an *a priori physics informed* approach strongly relies on the knowledge of physical processes involved. In some configurations, this information is not always available or could be difficult to extract. Thus, it would be ideal to have a systematic procedure to automatically isolate such energy patterns without the need for any knowledge or assumptions both on flow and geometry. In addition and more interestingly, a formalisation of sparsification techniques in the context of machine learning can be considered as an extension of the X-AI (eXplainable Artificial Intelligence) framework. This relatively novel field aims to provide a solid ground to the machine learning techniques trying to define a systematic methodology to ”explain” the results of a machine learning algorithm (Roscher et al., 2020). These techniques have received particular attention in recent times to integrate machine learning algorithms into applications where a solid understanding of physics involved or an explanation of the decision taken is required for safe utilisation of the device. These requisites are needed for a large number of applications ranging from medical diagnosis devices to autonomous land, sea and flying vehicles.

In the present work two different sparsification approaches have been developed as, schematically, shown in figure 1.1. First, the *a posteriori* approach involves an a posteriori evaluation of energy interactions between a fixed and immutable set of temporal and spatial modes used to generated the Galerkin model. Conversely, in the *a priori* approach the sparsity is encoded in the basis functions themselves and the sparse system is obtained directly by Galerkin projection of the Navier-Stokes equations onto the new basis function. As shown in figure 1.1 the steps corresponding to data collection and generation of the POD are in common for both approaches. As it will be shown later on, the choice of the POD as the starting point is not the only possible choice, even if it is recommended for its properties of orthogonality and energy optimality.

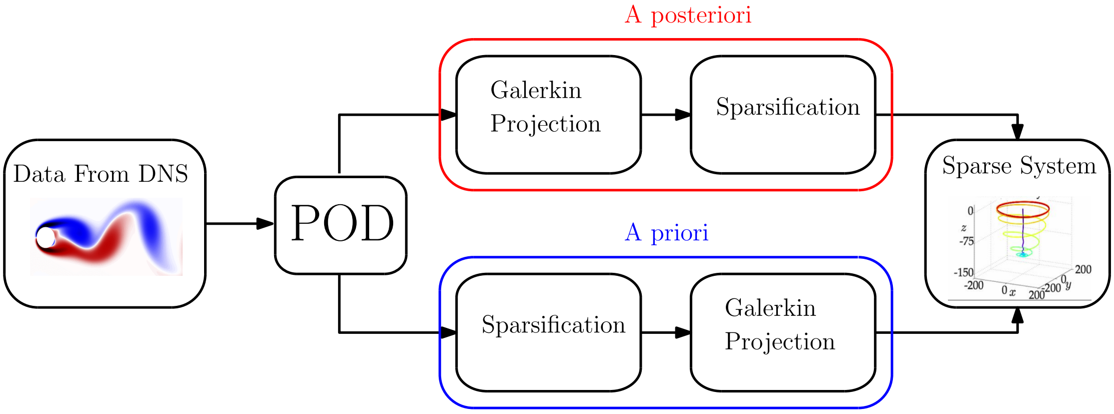

FIGURE 1.1: Schematic comparison between *a posteriori* sparsification in the red frame and *a priori* sparsification in the blue frame. Figures from Brunton et al. (2016)

1.2.1 *A posteriori* Sparsification

This technique, belonging to the class of the supervised learning algorithms, aims to identify the terms of the Galerkin model associated with relevant triadic interactions. The procedure can be sketched as displayed in figure 1.1 following the red branch. The first step consists in collecting the data and perform modal decomposition, we chose POD for later convenience but, as it will be shown later on, other choices are possible. The *a posteriori* framework, displayed inside the red frame in figure 1.1, consists in generating the reduced-order model first and subsequently perform sparsification on the originally dense models. The cornerstone of this approach is l_1-regularised regression (Tibshirani, 2013), widely used in the statistical community to extract a parsimonious representation of complex datasets containing a subset of predominant features. The non-differentiable nature of the l_1 regularisation allows transforming the interaction selection problem into a convex optimisation problem that can be efficiently solved, with a unique solution. Since no knowledge of the dynamics is utilised, the approach is fine-grained and relevant interactions are identified in a mode-by-mode fashion across the hierarchy of modes. However, since the sparsification procedure implies a modification of the system coefficients while keeping unchanged the spatial modes from which the coefficients are originated. As observed by Balajewicz et al. (2013) this could make the link between energy flows and modal structures questionable.

1.2.2 *A priori* Sparsification

The *a priori* sparsification technique aims to rigorously generate a set of modal structures with good energy representation properties such that the resulting interaction coefficient tensor is *a-priori* maximally sparse, without the need of subsequent post-processing or tuning of the model itself (see the blue frame in figure 1.1). For this task, we utilise POD as a starting point. Conversely, with respect to the *a posteriori* approach, here the sparsification is performed first rotating a set of POD modes within a POD subspace of larger dimension to obtain a new set of modes such that the resulting reduced-order model has a quadratic coefficient tensor with a sparse structure after performing Galerkin projection onto the subspace generated by these new modes. A particular example of this approach, where the subspace where the rotation is performed has the same dimension of the target subspace, is matrix diagonalisation. Here, through a change of basis, a representation of the same mathematical object where only the diagonal elements are different from zero is obtained [1].

1.3 Structure of the thesis

This thesis is organised following the progression of work done during the PhD and it is divided into four main sections. First, in the second chapter, the relevant works in the field are analysed to identify the main literature gaps and lay down a research plan and the objectives to be addressed during the research (explained in the third chapter). The fourth chapter is dedicated to a general discussion of the mathematical tools used for the generation of Galerkin-based reduced-order models and for evaluating energy flows within the Galerkin-model itself. Subsequently, in the second half of the chapter, the optimisation problems resulting from the two different sparsification approaches are discussed in details. Thus, the following two chapters, containing the body of the results, are organised into two macro sections, the sixth chapter is dedicated to the *a posteriori* approach (developed during the second year) while the seventh describe the *a priori* framework to which has been developed during the third and last year of research. The *a posteriori* methodology has been applied to three test cases of increasing physical and mathematical complexity. First, the Kuramoto-Sivashinsky equation has been chosen as a paradigmatic model containing multiscale dynamics. This model is particularly convenient since, due to the simple symmetry of the domain, the interactions can be conveniently described in Fourier space. This problem is particularly useful to understand how the sparsification algorithms act on the selection of the coefficients and on the reconstruction error of the sparse system. Second, we consider different two-dimensional problems, with particular attention on the lid-driven cavity flow. This test case is of particular interest since despite its relatively simplicity of the domain and of the boundary conditions, allowing an easy generation of the Galerkin model, it is possible to obtain solutions with arbitrarily complex physics. More specifically, just by increasing the Reynolds number, it is possible to obtain a set of different physical phenomena ranging from oscillating shear layer to a fully chaotic regime where the shear layer is broken down by the detachment of the vortexes generated in

---

[1] The code to diagonalise a matrix via subspace rotation is available in the following GitHub repository

the corners of the cavity. Lastly, once the framework has been proved effective, we consider a two-dimensional test case of more engineering interest, namely, the separated unsteady flow around a NACA 0012 profile at a high angle of attack. In addition, this test case is also used to test the computational framework in the case of non-trivial computational meshes.

The second result section is dedicated to the *a priori* framework. This novel methodology is fist applied to the lid-driven cavity flow with the same flow conditions ($Re = 2 \times 10^4$) analysed for the *a posteriori* approach. Similarly to what was done for the *a posteriori* sparsification framework, we have first tested the *a priori* approach on the cavity flow at the same Reynolds number as considered for the *a posteriori* methodology. The main objective of the cavity flow analysis is to validate the *a priori* sparsification methodology on a canonical test case. More interestingly, we will compare how these two different methodologies perform on the same flow configuration. Lastly, we applied the *a priori* sparsification technique to a more complex test case, namely, the separated flow over a backward-facing step. This configuration is a challenging test case for reduced-order modelling and its sparsification, mainly for two reasons. First, the open inlet and outlet boundary conditions introduce extra pressure term that needs to be modelled. Second, from the sparsification point of view, the strong convective nature of this flow generates a particular structure in the triadic interactions. It will be shown that the *a priori* formulation is able to capture this structure as opposed to the *a posteriori* one.

Lastly, in the closing chapter, we will compare the performance of both approaches on the common test case, the cavity flow, discussing the analogies and the relative strengths and weaknesses of both approaches. Finally, the best practices and how to choose the most suitable sparsification approach according to the problem considered is discussed. In the final section, possible future research avenues and practical applications will be debated.

Chapter 2

Literature Review

This chapter examines the relevant literature laying the groundwork for the present research. First, the advantages of model order reduction for describing high dimensional multiscale systems are discussed. Then, we introduce the state of the art of modal decomposition techniques and show how different formulations highlight different physical features of the flow. Second, energy flow analysis is introduced discussing the major results obtained for homogeneous turbulence and some more recent findings for flows evolving in more complex geometries. Third, some aspects of the large literature related to reduced order modelling are briefly discussed, explaining advantages and downsides of such an approach. Further in the chapter, we discuss some more recent ideas proposed to generate simplified models of turbulent flows such as, sparse regression methods for dynamical systems and *a priori physics informed* selection techniques for reduced-order models. Lastly, we discuss the general idea behind the l_1 based regularisation and why it works for model sparsification.

2.1 Analysis of coherent structures via modal decomposition

The multi-scale nature of turbulence makes a complete dynamical description of all features a demanding computational task, in addition and more importantly, this physical and mathematical complexity makes the interpretation of the solution difficult. A promising approach, receiving significant attention in recent years, is Reduced Order Modelling (ROM) or Model Order Reduction (MOR). This set of techniques leverages the idea of rewriting the original mathematical model (Navier-Stokes equations) on a new basis better representing the flow physics. Generally speaking, we can refer to modal decomposition as a set of techniques to systematically extract the most relevant flow features of the flows (Taira et al., 2017). The outputs are called basis functions and consist of pairs of spatial and temporal modes, generally accompanied by a characteristic value that can represent the energy, the frequency or the growth rate of the mode itself, according to the decomposition chosen. Once the spatial modes are obtained they can be used as a new basis for the construction of the reduced-order model by projecting the original model (Navier Stokes equations in the case of fluid dynamics) onto the subspace generated by a subset of the new modes, this procedure is called Galerkin-projection. Galerkin-projection results into a set of ordinary differential equations (ODEs) describing the evolution of the temporal modes and scaling down the computational cost of the solution with respect to the original Navier-Stokes. Generally, the new set of basis functions are obtained from geometrical considerations or from eigen-decomposition of operators associated with the dynamics of the system considered. According to the procedure chosen for their generations, basis functions can be classified into three main categories: mathematical, physical or data-based. Usually, the new basis functions are chosen to optimally represent physical aspects of the flow, such as energy reconstruction, spectral pureness, growth rate of the instabilities and so on. Therefore, performing a convenient modal decomposition of the flow field isolates fundamental flow structures associated with different physical phenomena. In addition, expanding the flow field onto an optimal modal basis generates more compact models. The advantages of such representation are twofold: it allows a deeper understanding of the flow physics, improving, at the same time, the computational efficiency of the model itself since fewer terms are considered in the model. We refer to Taira et al. (2017) for a complete overview of the advantages and drawbacks of the most popular modal decompositions.

Historically, model order reduction techniques leveraged mathematical modes. These consist of analytically defined functions and are derived from considerations on the domain geometry such as homogeneity or/and type of boundary conditions. The most popular and widely exploited example is Fourier decomposition, where the flow is decomposed as a finite sum of complex exponential functions. This approach naturally associates to each mode a well defined length scale making this approach preferable for the description of scale interactions in case of periodic boundary conditions. For example, Fourier modes are still the best choice for the analysis of homogeneous flows Pope (2001). A different approach consists in using physical modes. These are obtained from the analysis of a relevant operator derived from the physical model taken into consideration. As an example, in (Noack et al., 2011) the authors observed that it is possible to construct a complete set of basis functions suitable for describing the linear dynamics, from the linear global stability analysis of the Navier-Stokes operator. This technique was shown powerful since it leverages the vast tool-set related to the linear stability theory. However, it is inherently linear and difficult to extend to nonlinear flows, *i.e.* flows where different modes interact not weakly between each other. In this work we won't follow this approach since we are mainly interested in the description of statistically stationary fully developed turbulence. Since, in this flow regime, flow structures interacts non-linearly some authors (Taira et al., 2017) suggest the use of modal basis generated from data and/or anlysis of the nonlinear operator involved in the Navier-Stokes equations. The second example of a physical or operator based approach is proposed in Gómez et al. (2016). These authors proposed a different approach suitable to be applied to turbulent flows called resolvent analysis. The resolvent analysis aims to identify flow structures by using amplification analysis of the Navier-Stokes equations in the frequency domain. This decomposition allows expressing the velocity field as a sum of dynamically significant modes. These modes have the interesting property of being non-empirical and scalable with the Reynolds number, making the approach more general with respect to a purely data-based one. This technique has been successfully used to obtain a low-order representation of a turbulent flow in a channel and of the lid-driven three-dimensional flow inside a span-wise periodic cavity in Moarref et al. (2014) and Gómez et al. (2016). Interestingly, for these two configurations, it was shown that resolvent analysis acts in a very selective way and therefore it is possible to reproduce the dominant features of the flow with a low dimensional approximation. Finally, another interesting application of this decomposition has been used in Gómez et al. (2014). Here, starting from the observation of the presence of a sparsity patterns in the frequency domain, the authors proved the existence of a link between amplification and energy of the most energetically relevant flow structures identified by the resolvent analysis. Results suggest that the ability to successfully identify sparsity patterns in the frequency domain could be the key enabler to develop efficient reduced-order models based on selecting only the most energetically relevant modes.

All the aforementioned approaches are limited only to certain types of boundary conditions (see Fourier modes and Chebyshev polynomials) or to the knowledge of the mathematical model regulating the dynamics of the system (see resolvent analysis Gómez et al. (2016)). To fill this gap and extend model order reduction to a larger variety of problems, data-based modal decompositions have been introduced. This set of techniques, needing only data as input, has the advantage of not requiring any *a priori* assumption on the underlying dynamics and/or about the symmetries of the domain. The main drawback is the difficulty to generalise the results obtained for a certain flow configuration to another since, as the flow conditions changes, even slightly, a new modal decomposition has to be generated from scratch. Historically, the first data-based technique we mention is the Proper Orthogonal Decomposition, introduced by (Tennekes and Lumley, 1972; Holmes et al., 1997; Lumley, 2007). This technique has been widely used to analyse coherent structures in turbulent flows evolving in geometries of increasing complexity ranging from flat-plate boundary layer in Rempfer and Fasel (1994a) to stalled aerofoil in Kitsios et al. (2011). The main advantage of this technique is that the POD modes are optimal in the sense of capturing the maximum amount of variance in the data for each mode. For fluid applications this means that, defined a certain amount of kinetic energy to be reconstructed, the POD can recover it with the lowest possible number of modes, providing the optimal compression of the dataset in terms of the mean quadratic error. Energetic optimality, resulting in the

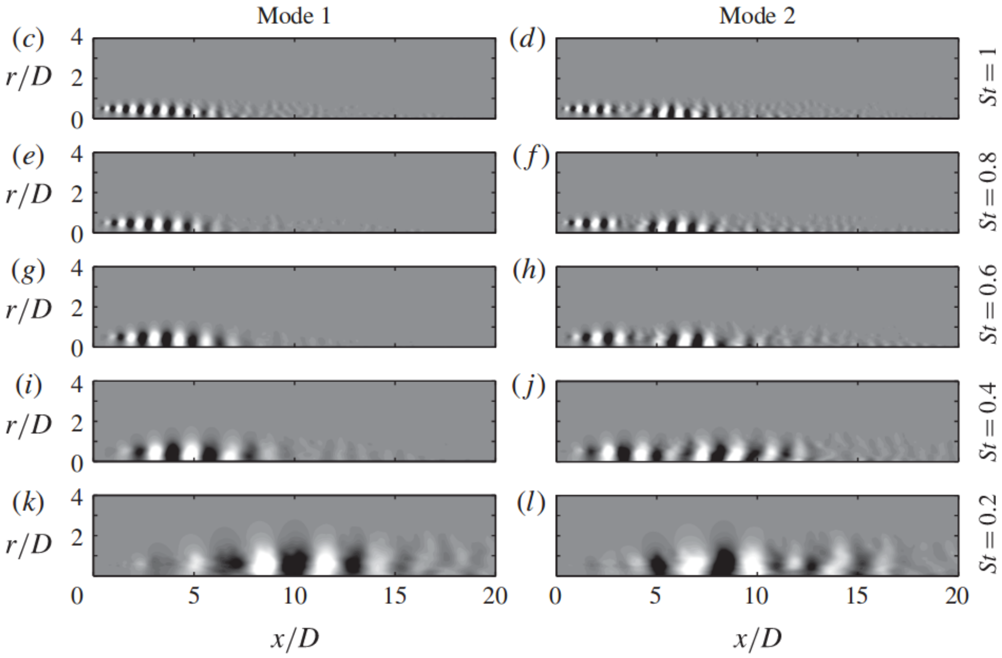

FIGURE 2.1: Spectral POD (SPOD) modes for turbulent jet. (c-l) Pressure field of the first (left column) and the second (right column) mode evolving at five different frequencies. Figure from Towne et al. (2018).

orthogonality between the modes both in space and time, makes this basis functions particularly suitable for generating compact and economic representations of the flow field and it has been widely utilised for the generation of reduced-order models. The main drawback is that the underlying mathematical procedure generates modes oscillating in time with spectrally complex time dependence, *i.e.* each mode contains an undefined blend of different temporal and spatial scales. As a consequence excluding particular flow configurations generating periodic flows, such that the vortex shedding around a cylinder (see Brunton et al. (2016)), it is generally not possible to ensure that a given POD mode contains only a predefined number of frequencies. As a result, every temporal POD mode contains an undefined blend of different frequencies associated with the different scales present in the system. This aspect is thoroughly explained in Towne et al. (2018) observing that this characteristic could become problematic when it comes to separate and analyse flow structures with similar energy content but occurring at different spatial scales. As a result, all these different spatial structures are lumped together in the same POD mode making the interpretation of the flow physics associated with each phenomenon challenging.

To overcome this issue POD in frequency domain has been developed. This approach generates a set of basis functions such that each mode contains a single temporal frequency but with the modes still ranked in terms of the energy. This decomposition is known as Spectral Proper Orthogonal Decomposition (SPOD). In the available literature, two different implementations have been so far proposed. The first approach by Towne et al. (2018) shares similarities with the original formulation proposed by Lumley in Lumley (1970). This procedure involves the decomposition of the spectral density tensor and leads to modes that contain only a single frequency. The authors made a connection between the classical space only POD and the SPOD showing that each POD mode is made up of several SPOD modes, each of them oscillating at a different frequency. From a physical standpoint, this means that each coherent structure, represented by a POD mode, is made up by the contribution of many coherent structures at different frequencies. This idea is shown in figure 2.1, where the first two SPOD modes for a turbulent jet at $Mach = 0.4$ are displayed. It can be observed that the first and most energetic SPOD mode (first column of figure 2.1) is composed by several flow structures oscillating at a prescribed frequency and then associated to a specific scale of the flow. The same observation can be done for the second SPOD mode reported in the second column of figure 2.1 and so on for the least energetic modes. A different

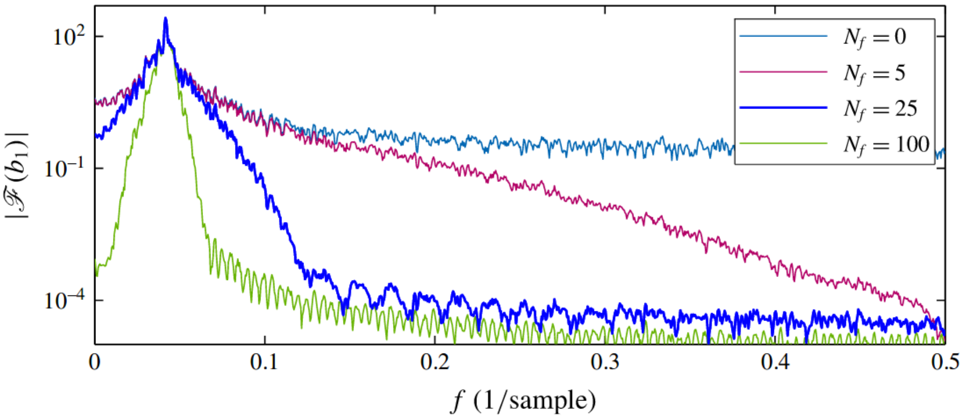

FIGURE 2.2: Magnitude of the frequency spectrum of the first SPOD mode, obtained with Sieber's approach, for three different width of the filter N_f. Increasing the width of the filter generates modes with a more pure spectral content. Figure from Sieber et al. (2016).

approach has been proposed by Sieber et al. (2016) where the authors developed an alternative algorithm to separate scales given the temporal correlation matrix. The main motivation of the authors is to bridge the gap between the classical energy ranked POD, the spectrally pure Discrete Fourier Transform (DFT) and the Dynamic Mode Decomposition (DMD) Schmid (2010), known to be challenging to use for turbulent realisation (see Taira et al. (2016)). This work was motivated by the observation that for highly turbulent flows coherent structures and intermittency phenomena could occur at similar frequencies making difficult to clearly separate the different phenomena. This makes necessary a large number of measurements and averaging procedures to obtain satisfactory results. The main idea proposed by the authors in Sieber et al. (2016) is to filter the temporal correlation matrix to force the resulting temporal modes towards a more clean temporal dynamics, *i.e* each mode should contain a narrow band of frequencies. The effects of the filter on the magnitude of the Fourier transform of the first SPOD mode of a turbulent jet are shown in figure 2.2. It can be observed that, increasing the width of the filter, the range of relevant frequencies contained in the temporal coefficient becomes narrower. Therefore, depending on the filter strength it is possible to range from the energetically optimal POD to the spectrally pure DFT. This ability to generate modes selective with respect the energy content comes at the cost of losing energetic optimality. This means that more SPOD modes are needed to reconstruct the same amount of energy, with respect to the POD. Moreover, SPOD modes are no longer orthogonal in space leading to some complexity when these basis functions are used to perform Galerkin-projection (see Taira et al. (2016); Rempfer and Fasel (1994a). Consequently, Galerkin models generated with SPOD modes will contain additional terms due to the inner product between different modes being non-zero.

Generally, for homogeneous isotropic turbulence, Fourier modes provide an optimal representation both in terms of frequency and energy content (Brasseur and Wei, 1994; Laval et al., 1999). For flows in complex geometries modes identified from data with Proper Orthogonal Decomposition (POD) and its spectral formulations are usually preferred (Schmid, 2010; Towne et al., 2018). In addition, we feel to mention a hybrid approach that aims to bridge POD and SPOD that has been recently proposed by Mendez et al. (2019). This algorithm, named multiscale POD (mPOD), leverages the idea of partitioning the spectrum of the correlation matrix, to produce modes that are energy ranked but containing only a narrow band of selected temporal frequencies.

However, these studies have utilised state-of-the-art modal decomposition methods that are inherently targeting optimal data representation. Therefore, these are not seeking for an optimal description of physical features such as sparsity of the interactions, nor include any information regarding the nonlinear structure of the governing equations that produce such interactions. Work in this direction has been recently done by Schmidt (2020) proposing a fresh approach aiming to generate modes that optimally represent flow structures associated with triadic interactions. One of the objective of this work is to widen the available literature by developing a modal decomposition with a maximally sparse triadic interaction tensor leveraging information about the mathematical structure of the equations.

2.2 Energy flow analysis

Energy transfer in turbulent flows has always been a topic of paramount importance to gain a better understanding of the physics of the process. One of the targets of this work is to develop a systematic procedure to select the most relevant features of the dynamics of the coherent structures in turbulent flows. Therefore, energy flow analysis is then a step of primary importance to understand how to leverage the physics of turbulent realisations to achieve our goal. Moreover, energy analysis is a fundamental step to check the physical preservation of the original properties once the model has been reconstructed. In this section, we will focus firstly on the results obtained in the framework of the modal decomposition of turbulent flows. Subsequently, we analyse some results obtained with different modal decompositons in different geometries.

2.2.1 Modal representation of the flow

The standard and well-tested approach to perform energy analysis of a turbulent flow makes use of Fourier modes. This mathematical description allows a very elegant representation of the spatial scales and offers a deep insight into the energy flow patterns between different structures. However, it remains confined to the cases of homogeneous turbulence. This limitation can be overcome by leveraging a data-based decomposition such as resolvent analysis or POD and its variants. Once the flow has been expanded, the evolution of the turbulent kinetic energy of the i^{th} mode can be defined as $e_i = \frac{1}{2}a_i^2(t)$, leading to the expression describing its temporal evolution:

$$\frac{de_i}{dt} = C_i a_i + L_{ij} a_i a_j + Q_{ijk} a_i a_j a_k + P_i(p, \phi_i, a_i) \tag{2.1}$$

Broadly speaking, the terms involved in equation (2.1) represent the work done by the mean field, the effect of the viscosity, the non linear interaction between the fluctuation and the work done by the pressure field. In this work we are mainly interested in the study of the nonlinear transfer term $Q_{ijk}a_i a_j a_k$. A more complete discussion of the terms involved in equation (2.1) can be found in section 4.1 and in Noack et al. (2011). Equation (2.1) is valid regardless the modal decomposition used to perform Galerkin projection, different basis functions affect the structure of the matrix coefficients.

In the early nineteens, some authors performed an interesting study reported in (Rempfer and Fasel, 1994a). They aimed to find a connection between the energy flows and the dynamical behaviour of coherent structures, defined as the POD modes of the flow, in a transitional boundary layer. The energy analysis suggests that the classical picture of the energy cascade adopted for homogeneous turbulence is not suitable for describing the energy fluxes between coherent structures in the considered test case. More specifically, in the classical energy cascade framework, the energy goes from the largest scales of the flow to the small ones and it is transferred by the interactions between structures of comparable size. On the contrary, they observed that energy is not just passed on from a structure to the nearest to the viscous scales, but rather, all the structures dissipate turbulent kinetic energy received from larger structures in parallel. It is arguable that this behaviour is mainly due to the presence of the solid wall that interacts with the flow. An important result is obtained temporal averaging the non linear transfer term $N_{ijk} = Q_{ijk}\overline{a_i a_j a_k}$. The authors have shown that the interactions' pattern between coherent structures is by no means arbitrary but, on the other hand, has a precise structure. Crucially, it can be observed that the most relevant interactions are the one grouped

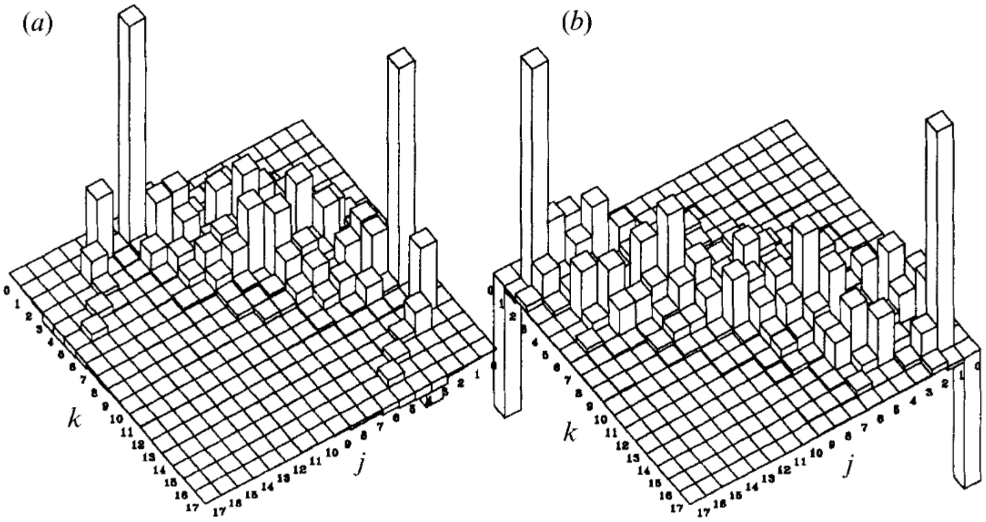

FIGURE 2.3: Histogram of the average energy flow N_{ijk} for $i = 12$ and $i = 17$ in panels (a) and (b), respectively. It can be observed that the most energetic interactions are clustered along the line defined by the equation $i = j + k$ and $i = |j - k|$. Figure from Rempfer and Fasel (1994a).

around the lines $i \sim j + k$ and $i \sim |j - k|$ in modal space, with i, j, k three different modal indexes. This pattern is plotted in figure 2.5 where the most energetic triadic interactions, defined as $Q_{ijk}\overline{a_i a_j a_k}$, are reported for two different values of i. This result is a direct consequence of the relation between triadic interaction in Fourier space. The only difference between this case and the spectral model is that for the present case this relation is only approximately satisfied by the POD modes. In addition, the authors in Rempfer and Fasel (1994a) suggest that it is possible to make a connection between the spatial structures and the indexing of the modes. It appears that the spatial size of a given mode is connected with the value of the index of the mode itself *i.e.* the most energetic modes associated with a low value of the modal index are also associated with the larger structures while the less energetic modes are associated with the small spatial structures. Therefore, the property of the POD decomposition of ranking the modes according to the energy content is reflected indirectly into an approximate rank also in terms of the spatial scales. More recently, in Towne et al. (2018) the authors observed that the separation between scales is a consequence of ranking the modes according to their energy content. However, it is important to point out that this property is not imposed by the mathematical formulation of the decomposition by is an indirect consequence of the energetic optimality. A similar analysis has been performed by Couplet et al. (2003). The authors performed a quantitative analysis of energy flows involved in the dynamics of a separated flow configuration past a backward-facing step. To understand the relative importance of the interactions between coherent structures they define the influence of the j-th mode on the i-th as,

$$\Pi_{ij} = \sum_{k=1}^{j} Q_{ijk}\overline{a_i a_j a_k} \tag{2.2}$$

Figure 2.4-(a) shows Π_{ij} on a logarithmic scale. Firstly, it can be observed that, in this case, as opposed to the Fourier representation all triadic interactions between the triplet of modes i, j, k, not only the ones satisfying the relation $i + j + k = 0$, contribute to the energy transfer. The energetic interactions occur in a local fashion since, as shown in figure 2.4-(b), the energy transfer between two modes i and j becomes more and more negligible as the difference between the value of the indexes $|i - j|$ increases. Lastly, it was observed that the energy transfer behaves qualitatively as observed for the Fourier decomposition *i.e.* the mean energy flow is directed from the large scales to the small ones of the spectrum via a local energy cascade mechanism. In addition, these authors suggested a physical explanation similar to the one provided by Rempfer and Fasel (1994a) noticing that the energy ranking of the modes is reflected on their spatial rank. This implies that for the higher indexes the POD modes converge towards the Fourier modes. This behaviour is qualitatively shown in figure 2.5 where iso-surfaces of the Q-criterion are displayed for four different modes in the first, mid and rear part of the spectrum respectively. It can be observed that as the value of the mode's index increases, the mode is associated to smaller and smaller structures and becomes qualitatively more similar to Fourier modes. Results obtained in Rempfer and Fasel (1994a); Couplet et al. (2003) have shown the complexity involved in the study of the energy transfer for complex flow configurations. However, it is worth underlining how for both configurations analysed only a subset of all the possible energetic interaction

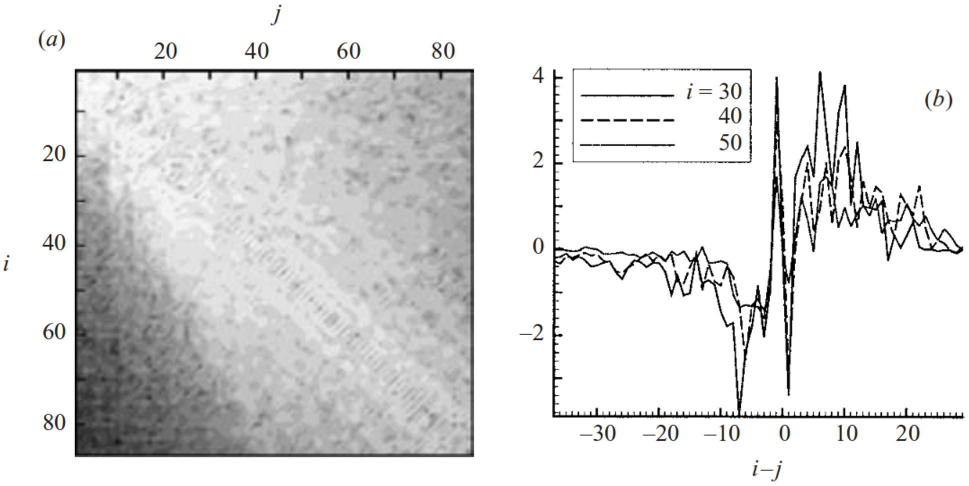

FIGURE 2.4: Panel (a) shows the map of $log(\Pi(i|j))$, white and black areas correspond lo large positive and negative values of $\Pi(i|j)$, respectively. Panel (b) shows horizontal cuts of $\Pi(i|j)$ as a function of $i - j$ for three different i. Figure from Couplet et al. (2003).

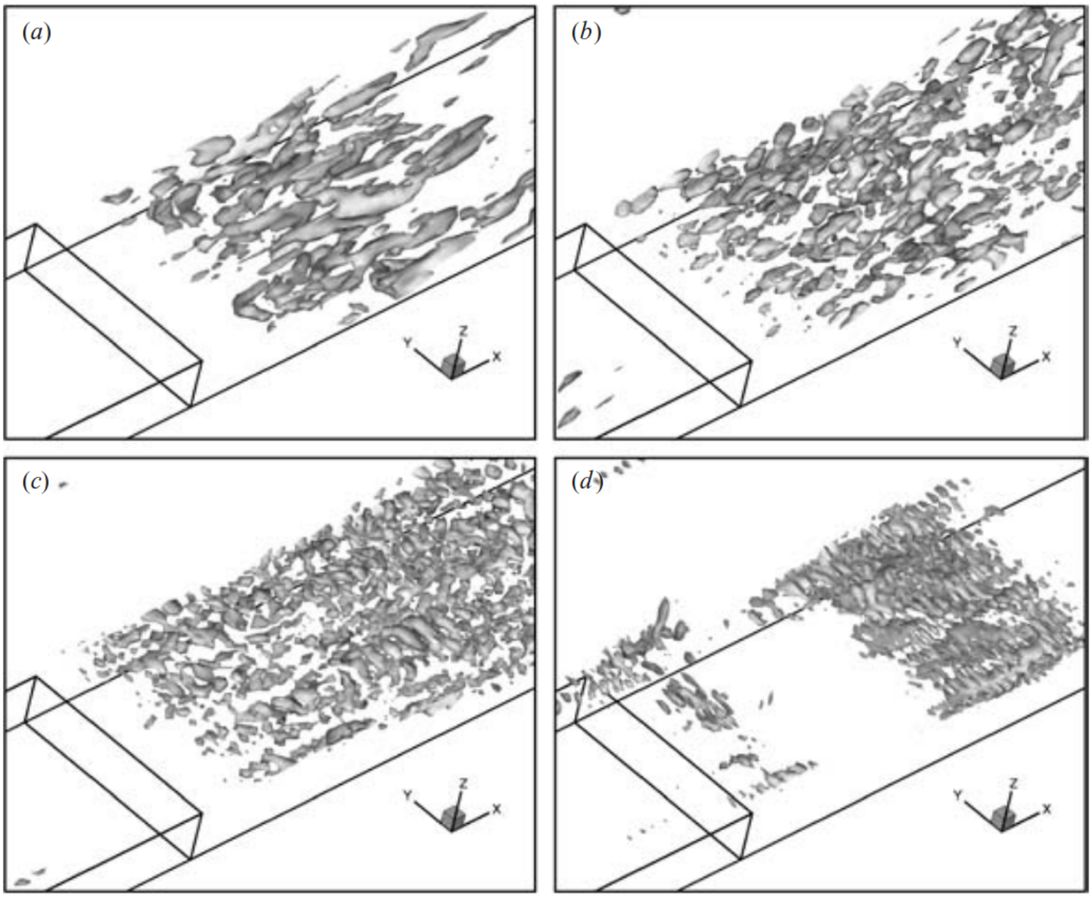

FIGURE 2.5: Iso-surfaces of the Q-criterion for the following POD modes: a) ϕ_1 b) ϕ_{20} c) ϕ_{42} d) ϕ_{87}. Figure from Couplet et al. (2003).

seems to play a relevant role in modal space. This is an important result confirming that generally speaking the flow considered are suitable for sparsification.

2.2.2 Homogeneous turbulence

Scales interactions in homogeneous turbulence is a deeply studied topic. Its importance lays in the fact that in the established picture of turbulence the energy transfer originated by the nonlinear term is the principal process that affects the evolution of turbulent flows. Despite the effort endured during the years by several researchers, the elementary processes involved still remain elusive. However, it is generally accepted that, at least for homogeneous flows, the energy transfer occurs in a local fashion *i.e.* the energy is transferred from the large scales of motion all the way down to the dissipative small scales via interactions between flow structures of similar size. This feature has been proven indirectly in Tennekes and Lumley (1972) where the authors underlined how the inertial subrange in the Kolmogorov spectrum is a direct consequence of the local interaction between triads. In Brasseur and Wei (1994), using data from numerical simulations, the authors observed that there is not appreciable energy transfer between wavenumbers widely separated in the spectrum, *i.e.* representing different spatial scales. Similar conclusions have been drawn in Domaradzki and Adams (2002) where the authors showed that the contribution to the energy budget of a certain scale comes mainly from scales of comparable size and therefore can be depicted as local. It is important to underline that all the works previously mentioned have been performed for fairly low Reynolds numbers and cannot be rigorously extended to higher Reynolds number flows. However, in works such as Rowley et al. (2009); Gómez et al. (2014); Brasseur and Wei (1994) it has been shown that the same mechanism holds, at least from a qualitative point of view, also for flows at higher Reynolds number.

In (Tennekes and Lumley, 1972; Brasseur and Wei, 1994) the authors suggest that in the context of triadic interactions is crucial to distinguish between two different concepts, *i.e.* energy transfer and triadic interactions. The energy transfer involves two different scales of motion. It can be either local if the two scales are similar or non local if it occurs between scales of different sizes. Triadic interactions represent the interaction between three different scales of motion. It can be local if all three scales are similar and non local if one or two of the three scales are bigger of smaller with respect to the other two. In the context of the triadic interactions, the concept of local or nonlocal energy transfer is complex. In fact, if local triad interaction always implies local energy transfer the opposite is not always true. It is possible to have local energy transfer caused by nonlocal triads if the energy is exchanged between two large wavenumbers in the triad without affecting the small wavenumber. Figure 2.6-(a) shows a local triadic interaction that occurs between modes with legs of similar size. While panel-(b) shows a distant interaction involving modes widely separated in the spectrum. In panel (b) is also represented the situation of local energy transfer between distant triads. An interesting result was found in Domaradzki and Rogallo (1990) where the authors analysed the transfer term involved in the energy equation. They found out that the energy transfer in homogeneous, isotropic turbulence is mainly local occurring between similar scales of motion. However, the triads interactions responsible for this process could

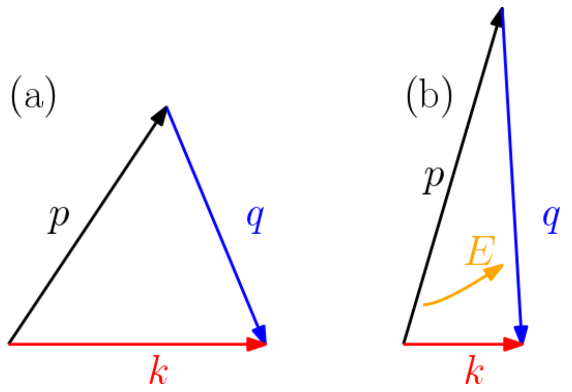

FIGURE 2.6: Schematic qualitative example of : local triadic interaction in panel (a) and distant triadic interaction with local energy transfer in panel (b).

be non local occurring between separated scales of motions. These results have been obtained with low Reynolds numbers turbulence but are likely to be extended to higher Reynolds numbers as well. However, from a global point of view, local energy transfer seems to dominate the dynamics of turbulent flows. The conclusions of Domaradzki and Rogallo (1990) are that, although they are unable to give a satisfactory description of the role played by the bigger scales in the energy transfer, the local energy transfer plays a predominant role in the energy transfer in particular in the vicinity of the spectral cutoff. The major takeaway of these works is that if this dynamics could be proved true it could indicate the possibility to simulate the biggest scales independently or more likely simulate the dynamics of certain scales taking into account only the scales of commensurate size.

Lastly, it is worth mentioning an analogous work in the field of two-dimensional turbulence performed by Laval et al. (1999). The authors performed the energy analysis of homogeneous two-dimensional turbulence showing that for the two dimensional case the energy interactions are non locals and only a small subset of interactions between large and small scales is actively involved in the energy transfer. In addition, these authors showed that it is possible to generate physically meaningful reduced-order models just by manually removing the less energetically relevant interactions. Results show that these manually pruned models reproduce faithfully the original physical features of the flow such as power spectra and amplitude fluctuation.

2.2.3 Energy interaction in different approaches

To conclude the overview on the energy analysis of reduced order models of turbulent flows it is worth mentioning a slightly different approach exploited by Taira et al. (2016) Nair and Taira (2015). The authors described the dynamics of two-dimensional turbulent flow as a network of interacting vortexes. The vortex network representation allows them to sparsify the system exploiting the well know techniques of network analysis Newman (2004) to identify the relevant interactions within the network. They show how to construct a sparse model of the network preserving the physical properties of the original one. This is an important result for our research since it is a further confirmation that although the equations describing a turbulent flow are dense they produce solutions containing an underlying sparse structure. In other words for the case of turbulent flows, it is possible to find a good sparse approximation of the original model.

2.3 Applications of ROMs

Reduced-order modelling is a powerful set of techniques allowing the generation of computationally cheap representations of large multiscale systems resulting from the discretisation of the PDEs. Hence, the affordable computational cost of the ROMs makes them an essential component to perform real-time simulations and control of complex PDEs systems. In particular, reduced order model techniques have been proposed by Ravindran (2000) Tonn et al. (2010) in the field of optimal control and optimal shape design Bui-Thanh et al. (2008) of flow devices. An interesting application in the field of flow structure interaction has been proposed by Lassila et al. (2012). In this case, they exploit the computational affordability of the ROM models to compute iteratively a large number of fluid solutions and reach the convergence of the fluid-structure interaction model. Because of the ability of the ROM models to condensate the most important features of the flow fields in just a few modes, they can also be used to perform more fundamental studies about the fundamental physics of the flow such as Alizard and Robinet (2011). These authors have successfully used the ROM approach to perform the stability analysis of a flat-plate boundary layer. Similarly, Kitsios et al. (2011) uses reduced order modelling approach to study the stability of the flow structures generated by the leading edge separation in NACA 0015 profile. Finally, it is worth mentioning that the reduced order-model techniques have been applied by Nguyen et al. (2014) to the framework of the reacting flow models.

Recently Gómez et al. (2016) proposed an interesting work developing a novel reduced order model for the three-dimensional unsteady lid-driven flow inside a span periodic cavity. This reduced-order model, obtained by resolvent analysis, requires only the mean flow and a minimal amount of information to be generated. The authors showed that the model is able to predict the velocity field with a good degree of accuracy using only a few number of modes. The main advantage of this approach is that resolvent analysis does not require any *a priori* assumptions about the nature of the non linearity of the flow and can be used for weakly non linear flow as well as for fully turbulent ones. Lastly, it is worth mentioning the work performed by Zhang and Wei (2017). These authors developed a reduced-order model of the vortex shedding behind a cylinder based on the dynamical mode decomposition. These authors introduce an approach to find the energetically dominant DMD modes enabling the generation of a set of basis functions suitable for the development of Galerkin based reduced order models while keeping the ability to isolate structures oscillating at a prescribed frequency. However, as it will be shown later on, the particular structure of the temporal modes generated by the DFT (coming in sin cos pairs) leads to some difficulties in mathematical conditioning of the optimisation problem involved in the sparsification procedure. This aspect will be discussed later on the section dedicated to the sparsification of the DFT-based reduced order models. Generally, unless in presence of particular flow configurations, for the generation of Galerkin based models, energy based decompositions have to be preferred.

2.3.1 POD based ROMs

One of the main goals of a Galerkin model is to construct a model containing the most important physical features of the original flow but distilled in the lowest number of degrees of freedom possible. The canonical approach consists in projecting the governing equations onto the subspace spanned by the modes generated by the proper orthogonal decomposition (POD). POD modes are computed optimising the average residuals with respect to the L_2 norm of the error between the original fluctuating field and the field obtained by modal expansion Holmes et al. (1997). This procedure, when applied to fluid dynamics, produces a set of spatial and temporal modes, $\boldsymbol{\phi}_i(\mathbf{x})$ and $a_i(t)$, maximising the amount of reconstructed turbulent kinetic energy for a given number of modes N. POD based reduced order models have been largely studied in literature and they have several advantages and just few drawbacks. In the following sections, we will briefly discuss the main problems of this approach and the solutions developed to tackle these issues.

2.3.2 Problematics of ROMs

Once the Galerkin projection of the Navier-Stokes equations is performed a coupled system of ODEs for the temporal coefficients a_i needs to be solved. Crucially, the computational cost of solving the ODEs system is much lower with respect to the one of solving the original PDEs system. However, it has been observed that the system of ODEs obtained simply by Galerkin projection produces predictions that are systematically way off the expected results. This outcome is expected and is a consequence of the implicit approximation done performing Galerkin projection onto a subspace of dimension usually much smaller than the original one. Therefore, some corrections need to be done to make the model suitable for practical applications.

Pressure Term

A compelling problem in model order reduction is related to the modelling of the pressure term arising from the Galerkin projection of the pressure gradient onto the subspace spanned by the POD modes. Generally, this term cannot be simply expressed as a combination of the POD modes and, except for particular configurations such as wall-bounded and periodic flows, it remains different from zero and unknown. To the author knowledge, several approaches have been proposed in the known literature. In Galletti et al. (2004) the pressure term is modelled as a linear combination of the POD modes $(\mathbf{u}, \nabla p) = C_{ij}a_j(t)$ and the coefficients of the matrix C_{ij} are found through an optimisation procedure aiming to match the predictions of the reduced-order model with the known dynamics of the flow. A different approach is proposed by Bergmann et al. (2009), they try to take into account of the pressure term formulating a pressure extend ROM Bergmann et al. (2008). In this new formulation, a new set of basis functions for the pressure field are introduced such that the pressure term can be directly evaluated through the pressure mode. Finally, an analysis of the problems involved by the pressure term and its modelling is provided by Noack et al. (2005). More specifically, the authors investigate the role of the pressure term in a shear layer configuration. They show how neglecting the contribution of the pressure term can lead to a non negligible amplitude errors even for a very low dimensional system. Most importantly, these errors cannot be compensated by increasing the number of modes taken into consideration.

Underpredicted dissipation

One of the most important features of POD-based ROMs is that, due to the energetic optimality of the POD, they capture a large amount of the kinetic energy of the flow in only a few modes. This means that in the reduced-order model only the bigger and most energetic structures need to be taken into consideration. This results in an aggressive truncation of the small and low energy scales where the viscous dissipation takes place. This is particularly true for high Reynolds number flows where a large number of modes is required to obtain a complete description of all the dynamically active scales. On the other hand, viscous dissipation is necessary to provide the correct physical description of the flow. This implies that ROMs built by only a few modes are not able to dissipate the correct amount of kinetic energy. This extra energy could lead both to inaccurate results or, in the worst cases, in numerical instabilities. Due to its direct practical implications, the stabilisation of aggressively truncated reduced-order models has been widely studied in literature and several approaches have been proposed to tackle the problem.

Bergmann et al. (2009) proposed two different approaches. In the first, called Residual Based Stabilisation, the goal is to approximate the fine scales of the flow with some adapted basis functions. The second method consists in updating the POD database on the fly during the simulation of the ROM itself. A more classical approach based on the concept of LES-like eddy viscosity model has been proposed by Östh et al. (2014). The authors present different formulations of increasing complexity to model the unresolved scales in the flow via the introduction of eddy viscosity. Then they tested these models for describing the flow around an Ahmed body configuration and showing good predictions capabilities of global features such as aerodynamics coefficients.

A different approach leverages the idea of automatically calibrating *a posteriori* the matrix coefficients of the ODEs system to introduce the correct amount of viscous dissipation. Recently, various approaches have been proposed to calibrate and improve the accuracy of the reduced-order models (Stankiewicz et al., 2011; Cordier et al., 2010; Amsallem and Farhat, 2008; Perez et al., 2017). They all share the idea of leveraging the knowledge of the true temporal dynamics of the flow to adjust the numerical value of the coefficients in the ODEs system through an optimisation procedure. A simple approach consists in performing least squares regression on the system coefficients. This approach can provide satisfactory results in the case of small ROMs and good quality data to train the algorithm. In real life configurations, as observed by (Cordier et al., 2010; Friedman et al., 2008), small and poor quality datasets are very common. This makes the use of a pure least squares regression technique, without any regularisation term, is strongly discouraged. Cordier and coworkers (Cordier et al., 2010; Galletti et al., 2004) proposed an approach based on a Tikhonov penalisation term to control the numerical values of the coefficients in the reduced-order model. This approach is a well established statistical learning technique used to regularise least squares regression problems where the database or feature matrix has poor numerical conditioning Hansen (1990); Friedman et al. (2008). Although the Tikhonov regularisation has been proved successful for relatively small models a possible drawback of this approach is that the L_2 penalisation preserves, in the tuned model, all the coefficients present in the original one. Since for POD-based Galerkin models the number of coefficients grows as the third power of the number of modes, the resulting Ridge problem might require a prohibitive number of snapshots to perform a meaningful regression without incurring in problems related to the overfitting. In addition, after the *a posteriori* calibration of the matrix coefficients, the physical link between spatial modes and ROMs coefficients is lost resulting in a questionable physical interpretation of the relative strength of the energetic interactions in the sparse model.

To solve this incoherence and preserve the physical link between the spatial modes and the ROMs coefficients a conceptually different solution, called subspace rotation technique, has been proposed by Balajewicz et al. (2013). The authors reproduce the right amount of dissipation in the model adding an additional constraint to the original proper orthogonal decomposition algorithm. The main idea of this procedure is to rotate the original set of basis functions into a new, more dissipative configuration and introducing the correct amount of viscous dissipation inside a low dimensional ROM model. The authors then have tested this approach over a high Reynolds number cavity flow proving its effectiveness. This approach is interesting because it is a very natural way to impose any constraint on the POD basis that is generated. In addition, subspace rotation resolves the problem of loss of coherence between the modal structure and new ODEs coefficients modified by the calibration since as the matrix coefficients are tuned the corresponding spatial modes are modified as well. This approach will be discussed in further detail later on and will be used in the *a priori* sparsification approach.

2.4 Model sparsification

The idea behind model sparsification relies on the observation that, although the mathematical model describing the dynamics of the flow is generally dense, energy flows between the hierarchy of structures follows preferential patterns. The idea behind model sparsification is to develop a procedure that automatically identifies these patterns without the need for any action by the user. This is beneficial since retaining in the model only the interactions that actively contribute to the dynamics allows to consider larger modal expansions including a wider range of flow scales improving the accuracy of the models and mitigating the effects of the truncation. In addition, a more compact representation of the triadic interactions can provide an easier interpretation of the dynamics involved helping to gain insight into the flow physics.

2.4.1 *Physics informed* sparsification

With the term *physics informed* sparsification we mean a procedure that exploits any previous knowledge about the physics of the flow or the geometry, such as symmetries, to simplify the problem itself. Recently, some authors exploited this idea to generate compacts model of turbulent flows. As an example in Thomas et al. (2014) the authors develop a Reduced Nonlinear Model (RNL) of a planar Couette flow modifying a DNS code retaining in the model only the terms that are supposed to contribute actively to the dynamics of the flow based on geometric observations. These authors observed that the restricted model can reproduce with a good degree of accuracy the mean velocity field and the longitudinal component of the Reynolds stresses. These results proved that the reduced model can capture the main physical features of self-sustaining turbulence. Most importantly the assumptions done make the model computationally and analytically tractable. This opens new possibilities of developing a deeper insight in understanding the dynamics of wall turbulence and developing new flow control strategies. A similar approach is presented by Gómez et al. (2014). They exploited the observations of Bourguignon et al. (2014) that near-wall turbulence presents a sparse representation in the Fourier space, *i.e.* only large sparse peaks are present in the spectrum. The authors suggest that these features are the direct consequence of the finite length periodic domain and more importantly that this property can be exploited to generate efficient reduced-order models based on selecting only the dominant frequencies among all the possibilities. Similarly, a *physics informed* approach have been used by Laval et al. (1999) exploiting the knowledge of energy fluxes in Fourier space to generate a minimal model for two-dimensional isotropic turbulence.

These approaches can provide an interesting insight on the physics. However, they still relies on *a priori* knowledge of the flow. Our aim is instead to develop a systematic procedure that can detect the most relevant features of reduced-order models of turbulent flows. A promising approach to this goal is the sparse regression technique that recast the problems of selecting the relevant features in a convex optimisation problem for which well-established routines exist.

2.4.2 Sparse regression of dynamical systems

Sparsity promoting techniques, alongside with machine learning, have been shown to be a powerful tool in many fields of engineering and science such as image recognition and signal analysis Roig et al. (2009); Blum and Langley (1997) . Only recently, these concepts have been extended to a more physics-related framework. More specifically, statistical learning techniques have found a promising application in the identification of dynamical systems from data. Sparsification techniques can be applied to the reconstruction of dynamical systems from data, as explained in figure 2.7, in four main steps. First, snapshots of the velocity or vorticity field are collected and the corresponding modal coefficients $a_i(t)$ are computed. Subsequently, assuming a mathematical structure of the system, the data are recast in a matrix form suitable to perform linear regression. Thus, the l_1 regression problem is solved to identify only the relevant terms for the temporal dynamics of the system and ignore the others, whose corresponding coefficient β_i will be set to zero. Mathematically this is done setting to zero the coefficients β_i whose absolute value goes below a certain threshold proportional to the value of the regularisation weight λ_i. Significant efforts in this direction have been reported in Brunton et al. (2016); Loiseau and Brunton (2018). These authors developed the SINDy (Sparse Identification of Nonlinear Dynamics) framework that exploits the sparsity promoting techniques such LASSO (Least Absolute Selection Shrinkage Operator) regression and subset selection to reconstruct nonlinear governing equations from noisy measurements. The SINDy methodology is schematically explained by the blocks of figure 2.7. This method has been successfully tested on different nonlinear

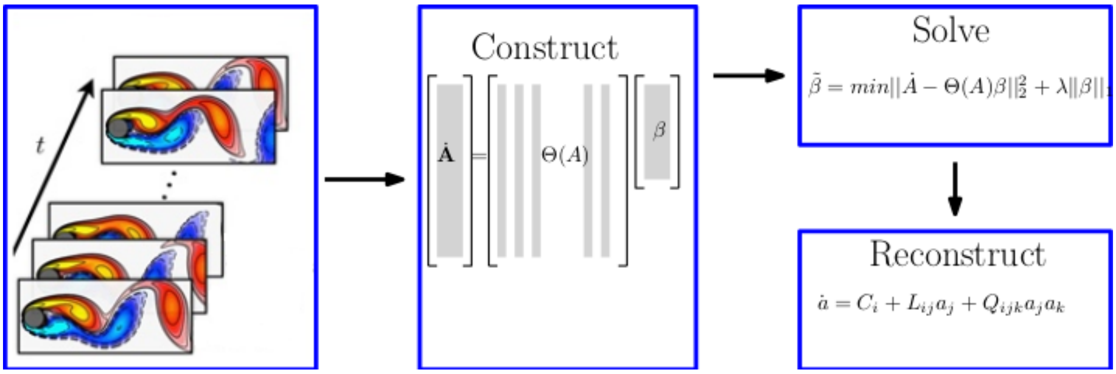

FIGURE 2.7: Algorithm to reconstruct dynamical systems from data. It consists in four steps shown in the blue frames in the figure above. 1) The data are collected from simulation or experiments. 2) the matrix of the time derivative $\dot{A}$ and the library $\Theta(A)$ are constructed. 3) the LASSO is performed. 4) the dynamical system is identified.

chaotic systems such as the Lorenz attractor and the POD reduced-order model of a laminar flow around a cylinder. The results reported in Brunton et al. (2016) show how the LASSO algorithm is able to identify the correct structure of the dynamical system such as limit cycles and attractors. These works are of key importance for the extension of the sparsification algorithm to a more physical framework, showing that the sparsification algorithm can reconstruct and identify the correct physics of the system. A further improvement of the SINDy algorithm has been done in Loiseau and Brunton (2018). In this work, the authors argued that the sparse regression technique is just a mathematical procedure performed on the dataset but the conservation of the physical properties of the original system is not automatically granted and must be ensured during the mathematical procedure. The authors then presented a way to enforce the energy conservation of the nonlinear term in the ROM representation of the Navier-Stokes system. To this goal, they formulated the LASSO problem as a convex minimisation problem with additional linear constraints on the system coefficients. Then the problem has been solved by the optimisation framework cvxpy Diamond and Boyd (2016). The energy conservation of the nonlinear term of the Navier-Stokes equations is an important aspect to preserve the numerical stability of the ROM and to ensure that the physics of the original flow will not be modified once the sparsification is terminated. Both of these issues are considered in Balajewicz et al. (2013) and Jovanović et al. (2014) where a mathematical reformulation of the POD and DMD is proposed to obtain a reduced-order model that describes the correct amount of dissipation and energy transfer between the scales and select only the relevant modes. More interestingly, this approach involves the modification of the basis functions therefore the link between the flow structure and system coefficients is not lost. A similar idea will be explored later on in this work when the *a priori* sparsification approach is discussed.

Although sharing a similar approach our perspective is different and closer to the work of Taira et al. (2016); Nair and Taira (2015). These authors described the dynamics of a two-dimensional turbulent flow as a network of vortexes and then applied the network sparsification techniques. These techniques allow describing the interactions between the node of the network through the adjacency matrix. Several mathematical techniques have been set up to sparsify the adjacency matrix of a network Newman (2004). The result is qualitatively reported in figure 2.8. In figure 2.8-(a) and (b) two different sparsity patterns for the adjacency matrix are plotted. The colour of the cells is the absolute value of the weight associated with interactions between two vortexes. It can be observed how the highest values near the diagonal represent the presence of a cluster of vortexes. Figure 2.8-(c) shows the eigenvalues of two different graphs with different sparsification ratios compared with the original ones. It can be observed that the numerical value of the eigenvalues is slightly affected by the sparsification procedure. This work showed how it is possible to extract a sparse structure from an inherently dense system without any *a priori* knowledge on the structure of the sparsity structure of the system itself.

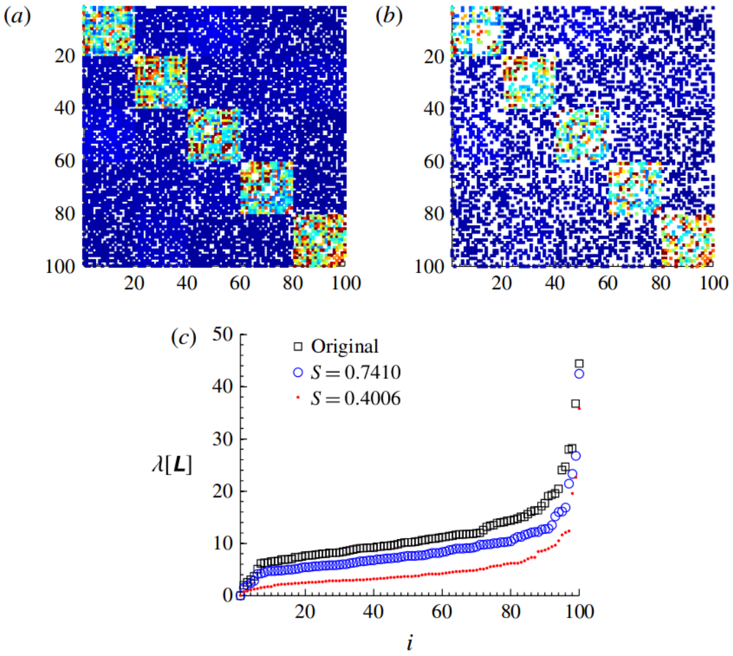

FIGURE 2.8: Panels (a) and (b) show the sparsity pattern of two different networks of vortexes. The colour represents the absolute value of the weight of the interaction described by that term, red and blue cells correspond to high and low values, respectively. Panel-(c) Eigenvalues of the spectra of the two sparsified system compared to the original one. Figure from Taira et al. (2016).

2.4.3 Automatic selection of relevant features

This work heavily relies on Machine and Statistical Learning techniques, in this section a brief overview of this set of techniques will be done. In recent years Machine Learning (ML) started to spread from purely Information technology applications towards more physics and engineering-oriented applications. Since large datasets have become readily available Machine Learning is a fundamental tool to identify patterns and relevant features inside datasets whose size would be overwhelming for a person to analyse. An extensive overview is done in Blum and Langley (1997). The authors explain at a conceptual and practical level the problem of ”relevance” in the data mining framework. The bottom line is that at a practical level we seek an algorithm that scales well with domains with many irrelevant features. More specifically, the number of training samples needed to reach a certain level of accuracy should grow slowly with respect to the number of features present in the dataset. A good example of a feature selection model is explained in Roig et al. (2009). A novel point of view to create a correspondence between two different images is presented. The key point is that the problem can be reformulated as the optimisation problem of a convex energy function representing the distance in a certain subspace between the reconstructed image and the target image. They also apply three different approaches to solve the minimisation problem showing the superior performances of the approach based on the convex quadratic programming.

The previous examples although very interesting have the common property to not have any clear structure in the original dataset. This is not the case of the reduced-order model of a flow that we know to have a polynomial dependency from the temporal amplitude of the modes. The particular structure of the data allows us to restrict the searching area by implementing a linear regression on a function database defined *a priori* by the user (see Friedman et al. (2008)). The linear regression framework is particularly appealing since they have a fairly simple mathematical formulation and implementation. In addition, these techniques can provide adequate accuracy and interpretability of the results in particular for models where the number of features is bigger than the number of training samples (situation likely to occur in fluid dynamics applications) they can provide even better results and clearer interpretability that more complex algorithms. A nice review about Machine Learning techniques applied to fluid mechanics is provided by Brunton et al. (2019) where the state of the art of machine learning techniques in fluid dynamics and an overview of the future challenges are presented.

2.4.4 Role of the regularisation in linear regression

Often times when complex datasets are used in a linear regression framework, the number of features is smaller than the number of available samples. As pointed out by (Roig et al., 2009; Cordier et al., 2010) in this scenario a simple least-squares approach will incur into overfitting making the results meaningless. A simple solution is including in the mathematical formulation of the problem an extra penalisation term in the form of some norm of the vector coefficients with the aim of regularising the solution making the optimisation problem well-posed. In addition and more interestingly, the mathematical form of the regularisation term is used to enhance different features of the dataset (Friedman et al., 2008). A good overview of the linear regression techniques

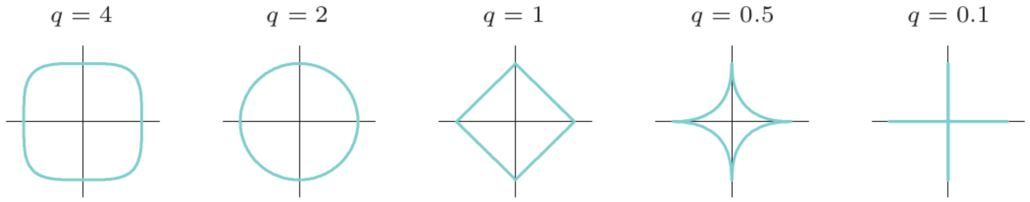

FIGURE 2.9: Iso-contours of $||\beta_i||_p$ for a constant value. The image shows decreasing values of q moving from left to rigth. Image taken from Friedman et al. (2008)

and data mining is provided by Roig et al. (2009). Generally speaking, it is possible to define for any real number $p > 0$ different norms as

$$||\beta_i||_p = (|\beta_0|^p + |\beta_1|^p + ... + |\beta_I|^p)^{1/p} \qquad i = 1, ..., I \tag{2.3}$$

where I is equal to the number of element of β_i. Two limiting cases of this definition are: the l_0 norm that only counts the number of elements of β_i different from zero, also known of carnality, defined as $||\beta_i||_0 = card(\beta_i)$ and the l_∞ norm identifying the largest element of β_i, defined as $||\beta_i||_\infty = max(|\beta_i|)$. The effect of different types of norms can be understood by visualising the contours of constant values of $||\beta_i||_p$ in the space of the optimisation variables β_i. Figure 2.9 shows the values of the contours for decreasing values of q from left to right. Starting from the very left we have q increasing leading eventually to the l_∞ norm. The regularisation obtained with this norm, called subset selection, generate a solution that takes into account only the largest element of β_i. The major drawback of this approach is that a feature is either kept or eliminated from the selected subset. Consequently, very different solutions can be obtained with a slight change in the training dataset and/or of the hyper-parameters of the regression. Moving right we have the case $q = 2$, this is the well known Ridge regression where the coefficients are shrunk to zero by a quadratic penalisation, although this regression is more stable and provides statistically better results with respect to the pure least-squares it still produces a dense result and then hardly interpretable. This kind of regularisation has been used by (Cordier et al., 2010) to tune the predictions of POD based reduced-order models. Moving further to the right we obtain the case with $q = 1$. This configuration correspond to the LASSO regression Tibshirani (1996). We refer to Friedman et al. (2008) for a more in-depth discussion of the functioning of the l_1 norm. A crucial aspect of the family of sparsity promoting techniques is what happens to the contours shown in figure 2.9 when crossing the the value $q = 1$. More specifically, the sparsity of the solution is promoted by the non-differentiability of the solution on the axis. This produces non-vanishing gradients near the axis and "pushes" the small coefficients to be exactly zero. This is equivalent to performing a iterative thresholding procedure on the coefficients β_i. This approach has been used by Brunton et al. (2016) for system identification. In addition, we refer to Friedman et al. (2008) for a more broad discussion about the relation between different values of q, type of thresholding and the resulting coefficients selection strategy. The sparsity is promoted by the presence of spikes or non differentiable points in the shape of the norm. Therefore every norm such that $q \leq 1$ is, potentially, a sparsity promoting norm. However only for the case $q = 1$ the objective function is convex, therefore efficiently solvable. Lastly, as the value of q decreases, we observe that the non differentiable points become sharper and at the same time the norm becomes non-convex. The limit configuration is the l_0 norm that simply counts the non-zero elements of the vector β_i. From a theoretical point of view, the l_0 norm would be the optimal one to promote sparsity in the solution. However, as observed by Jovanović et al. (2012, 2014), its non-convex nature makes the solution of the problem computationally hard and unpractical to solve for systems of already modest dimensions. For this reason, the best trade-off between sparsity and computational tractability is offered by the l_1 penalisation or LASSO regression. In addition, it can be shown that the l_1 norm is the best convex approximation of the "optimally" sparsity promoting l_0 norm (Ramirez et al., 2013). Crucially, the LASSO does not focus on a certain subset of coefficients but defines a continuous shrinking operation leading to solutions with coefficients that are exactly zero. This allows defining a whole family of sparse reduced order models with different accuracy sparsity trade-off. Due to its convex and sparsity promoting properties we will use the l_1 norm to construct the sparsity promoting operators that will be introduced in the methodology section. In addition, the l_1 based regression is simple to implement, mathematically robust and its results are easily interpretable. All these properties are important assets for a machine learning model.

2.5 Summary and open questions

The literature review has shown that several modal decomposition techniques exist, each of them aiming to better highlight different physical aspects of the flow. The choice of the most suitable technique is not straightforward and mainly depend upon the characteristics of the flow considered. Broadly speaking, modal decompositions can be divided into two macro-categories. First, energy-based approaches, ideal to obtain a compact representation of the flow, can isolate flow structures with large energy content. However, they lack the ability to separate events happening at different scales but having similar energy content, making the interpretation of certain flow configurations challenging. The other class of decompositions consists of frequency-based approaches aiming to decompose the flow into modal structure oscillating at a fixed frequency. This set of techniques works better in separating events happening at different flow scales and describing statistically stationary flows. However, they show some limitations in configurations with a more complex time behaviour such as high intermittency and transient flows. During the years hybrid approaches, aiming to join these two categories of modal decompositions, have been proposed. Arguably, the major limitation of the state of the art is that they fundamentally seek an optimal representation of the dataset without considering neither the physics nor the mathematical structure of the system considered. Only recently some work in this direction has been done by Schmidt (2020) formalising a modal decomposition optimally representing structures associated with triadic interactions.

Once a modal basis, to expand the flow has been chosen, a set of ODEs describing the evolution of the modal structure can be generated. Generally, these models have a dense structure, *i.e.* the evolution of one flow structure depends upon the complementary set of structures. However, several studies show that triadic interactions do not have the same importance but that energy flows according to preferential patterns. On the other hand, when Galerkin-based reduced-order models are generated this sparsity feature is generally lost.

Recently some work has been done to integrate statistical learning techniques into fluid dynamics. More specifically, l_1 sparse regression has been successfully employed to identify dynamical systems from data. Despite the efforts put into the understanding of these methodologies, it is still unclear how they interact with the conservation properties of the flow, such as energy conservation of the nonlinear term and temporal stability of the sparse reduced-order model. More specifically, if sparsification techniques can preserve the correct physics of the flow or if it has to be manually enforced. In addition, to the author's best knowledge, these techniques have been applied so far to canonical low dimensional test cases having already a sparse structure embedded in the mathematical structure of the model and no applications to complex, large Galerkin models are present in the literature. Some work in this direction has been done in the context of vortex networks describing the dynamics of a two-dimensional turbulent flow as a network of vortexes and then applied the network sparsification techniques. Interestingly, the authors showed how it is possible, starting from a fully connected network, to remove links between the vortexes affecting only marginally the properties of the network. This result shows that it might be possible to sparsify an initially dense model exploiting the sparse nature of the energy flows inside the flow.

Chapter 3

Objectives

The goal of the present research is the development of a systematic methodology to generate computationally cheap, but yet accurate, numerical models of turbulent flows. More specifically, one of the major issues in the mathematical description of complex flows is the intractable number of degrees of freedom needed to resolve all the scales of the flow. A possible solution consists in projecting the governing equation onto the subspace generated by a suitable modal decomposition. In recent years, POD modes have become a very popular approach to generate the low dimensional subspace. However, POD-based models are mathematically dense, *i.e* the matrices involved in the definition of the reduced-order model have entries different from zeros, leading to a high computational cost and difficult physical interpretation. The key idea proposed in this work consists in exploiting the physical observation that non all the flow structures interact with the same strength, this generates a subset of dominant energetic interactions within the model. To our goal, we aim to develop a data-driven methodology as means to identify relevant triadic interactions in Galerkin models of turbulent flows. As opposed to a mere truncation, where the modes describing less energetic flow structures are simply ignored, the philosophy of this approach consists in generating reduced order models from a large number of modes and consequently, that can describe a large range of spatial scales. The computational efficiency is then enhanced by eliminating the least relevant interactions within the dynamics of the single modes. As a result, the temporal evolution of a modal structure requires the evaluation of fewer terms with respect to the original model. The cornerstone of this approach is the l_1 based sparse regression. Its mathematical formulation consists in solving an optimisation problem composed by a least-squares part and an l_1 penalisation term on the matrices coefficients of the reduced-order model. Crucially, this mathematical formulation translates the problem of identifying the most relevant interactions into a well posed optimisation that can be efficiently solved leading to a unique solution. The first part of the research has focused on a supervised learning-based sparsification technique we will call *a posteriori* sparsification. The *a posteriori* sparsification procedure will be applied to large Galerkin reduced-order models of turbulent flows. The first aspect that will be analysed is the ability of the l_1 framework to sparsify initially dense models in accordance with the pattern of the most relevant triadic interactions. A second crucial aspect, in the *a posteriori* sparsification, is to understand whether it is possible to naturally embed this mathematical procedure into the fluid mechanics framework in such a way the conservation laws and the symmetries implicit in the Navier-Stokes equations are preserved by the sparsification procedure. In case of a negative answer, we need to understand how to automatically enforce these properties in the formulation of the regression. The second contribution of this research is more fundamental in nature. It concerns the relation between sparsity of energy interactions and the definition of the modal basis used to generate the Galerkin model. More specifically, the classical l_1 based regression, proposed as our first approach, involves an *a posteriori* evaluation of the energy interactions between an immutable set of modes, pruning weak interactions that do not contribute to the dynamics. The downside of this approach is that the system coefficients are modified by the regression since the minimisation of the error between the original and the sparse system involves fine-tuning of the nonzero model coefficients without modifying the basis function. Consequently, the strict link between spatial structures and projection-based coefficients is lost making the interpretation of the modal energy transfer questionable. Arguably, an elegant solution is to define a new set of modal structures such that the corresponding quadratic interaction tensor obtained by Galerkin projection onto the subspace generated by these new modes is sparse. The key idea is to obtain the new modal functions as a linear combination of a preexisting set of modes. This is done by formulating an optimisation problem in terms of a rotation matrix linking the new to the old basis functions. The properties of the new basis will be enforced by including additional constraint constraints on the mathematical structure of the model involved. Classically, modal decompositions are educed from energetic optimality considerations and no pieces of information about the physics nor the mathematical structure of the model are considered. Thus, this approach aims to widen the current literature proposing a methodology able to include information about the mathematical model considered directly inside the formulation of the problem. Due to the novelty of the *a priori* sparsification approach two main goals have been defined. First, provide a robust mathematical foundation of the *a priori* sparsification method. More specifically, it is of paramount importance to understand how our objective can be efficiently translated as an optimisation problem. In particular, due to the mathematical structure of the third order tensor, we expect the optimisation problem to be not convex. How this aspect impacts the computational cost and the uniqueness of the solution is an aspect that needs to be analysed. Second, this approach will be demonstrated on canonical flows of increasing complexity. Here the goal is to understand how the new basis function is able to capture energy transfer mechanisms efficiently, *i.e.* where the quadratic coefficient tensor, regulating triadic interactions, is as sparse as possible.

Chapter 4

Methodology

This section outlines the mathematical tools required to generate Galerkin based reduced-order models and analyse energy transfers between modal structures. Subsequently, it is shown how to recast Galerkin models in a suitable form to perform *a posteriori* and *a priori* sparsification. Finally, these two opposite methodologies are compared discussing the major advantages and drawbacks of both methodologies.

4.1 Model Order Reduction Techniques

We consider the evolution of an incompressible flow defined by the velocity components $\mathbf{u}(t, \mathbf{x}) = (u, v)$ in a domain Ω of arbitrary shape described by the Cartesian coordinates $\mathbf{x} = (x, y)$ and endowed by the standard inner product

$$(\mathbf{u}, \mathbf{v}) = \int_{\Omega} \mathbf{u} \cdot \mathbf{v} \mathrm{d}\Omega, \tag{4.1}$$

where $\mathbf{u}, \mathbf{v}$ are two elements of such space. The resulting $\mathcal{L}^2(\Omega)$ norm is denoted as $||\mathbf{u}||_2 = \sqrt{(\mathbf{u}, \mathbf{u})}$. The equations describing the evolution of the velocity field $\mathbf{u}$ in every point in the domain are

$$\partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p^* + \frac{1}{Re} \nabla^2 \mathbf{u} \tag{4.2a}$$

$$\nabla \cdot \mathbf{u} = 0, \tag{4.2b}$$

where ∂_t denotes the derivative with respect to time, ∇ the gradient operator with respect to the spatial directions $\mathbf{x}$ and p^* the reduced pressure. The Reynolds number is defined as $Re = UL/\nu$, where U,L are reference length and velocity and ν is the kinematic viscosity. Once obtained a solution of (4.2), using the time averaged (indicated with the overbar) velocity field $\bar{\mathbf{u}}(\mathbf{x})$ as base flow, and denoting by $\mathbf{u}'(t, \mathbf{x})$ the velocity fluctuations $\mathbf{u}(t,\mathbf{x}) - \bar{\mathbf{u}}(\mathbf{x})$, a N-dimensional expansion expressed by the ansatz

$$\mathbf{u}(t,\mathbf{x}) = \bar{\mathbf{u}}(\mathbf{x}) + \mathbf{u}'(t,\mathbf{x}) = \bar{\mathbf{u}}(\mathbf{x}) + \sum_{i=1}^{N} a_i(t)\boldsymbol{\phi}_i(\mathbf{x}), \tag{4.3}$$

is introduced to describe the space-time fluctuating velocity field, where $a_i(t)$ and $\boldsymbol{\phi}_i(\mathbf{x})$, $i = 1, \dots N$, with $\|\boldsymbol{\phi}_i(\mathbf{x})\| = 1$, are the temporal and spatial modes, respectively. These modes may be computed *a posteriori* from numerical or experimental data, *a priori* from a characteristic operator of the system (Taira et al., 2017) or from completeness considerations (Noack and Eckelmann, 1994). Reduced order models are then derived by projecting the residual of the governing equations (4.2) defined as

$$\mathbf{R}(\mathbf{u}) = \partial_t \mathbf{u} - \mathbf{u} \cdot \nabla \mathbf{u} + \nabla p^* - \frac{1}{Re}\nabla^2 \mathbf{u}, \tag{4.4}$$

onto the N-dimensional subspace spanned by (4.3) via the inner product defined in (4.1). This is done by imposing the error to be normal to the subspace utilised for the projection

$$(\boldsymbol{\phi}_i, \mathbf{R}(\mathbf{u})) = \int_\Omega \boldsymbol{\phi}_i \cdot \mathbf{R}(a_i \boldsymbol{\phi}_i) d\Omega = 0 \qquad i = 1, \dots, N. \tag{4.5}$$

Restricting our analysis to configurations where the boundaries are either no-slip walls or periodic, procedure (4.5) results in an autonomous system of coupled nonlinear ordinary differential equations (ODEs)

$$\sum_{j=1}^{N} M_{ij}\dot{a}_j(t) = \tilde{C}_i + \sum_{j=1}^{N} \tilde{L}_{ij} a_j(t) + \sum_{j=1}^{N}\sum_{k=1}^{N} \tilde{Q}_{ijk} a_j(t) a_k(t), \qquad i = 1 \dots, N, \tag{4.6}$$

defining the temporal evolution of the coefficients $a_i(t)$. For the case of different boundary conditions an extra term involving the scalar product $(\boldsymbol{\phi}_i, \nabla p^*)$ is present in (4.6).As pointed out by Noack et al. (2005) this term cannot be computed directly in the solution of (4.6) but needs to be modelled. In addition, as pointed out by (Noack et al., 2005) increasing the number of modes does not improve the performances of the model. Therefore, it's modelling is a crucial step in the development of reduced order models. In this work two different approaches are used. First, in the testing phase of the algorithm, domains with closed walls are considered where the pressure term is identically zero. Second, in the case of more complex configurations we did not include any pressure term but we leave up to the sparsification procedure the tuning of the coefficients. This approach is similar to the one successfully used by (Cordier et al., 2010).

Here, we report the mathematical definitions of the coefficient tensors $\tilde{\mathbf{C}}$,$\tilde{\mathbf{L}}$ and $\tilde{\mathbf{Q}}$ in system (4.6) whose entries are defined by the expressions:

$$\tilde{C}_i = (\boldsymbol{\phi}_i, \bar{\boldsymbol{u}} \cdot \nabla \bar{\boldsymbol{u}}) + \frac{1}{Re}(\boldsymbol{\phi}_i, \nabla^2 \bar{\boldsymbol{u}}), \qquad i = 1, \dots, N, \tag{4.7}$$

$$\tilde{L}_{ij} = \frac{1}{Re}(\boldsymbol{\phi}_i, \nabla^2\boldsymbol{\phi}_j) + (\boldsymbol{\phi}_i, \boldsymbol{\phi}_j \cdot \nabla\overline{\boldsymbol{u}}) + (\boldsymbol{\phi}_i, \overline{\boldsymbol{u}} \cdot \nabla\boldsymbol{\phi}_j), \qquad i, j = 1, \ldots, N, \tag{4.8}$$

$$\tilde{Q}_{ijk} = (\boldsymbol{\phi}_i, \boldsymbol{\phi}_j \cdot \nabla\boldsymbol{\phi}_k), \qquad i, j, k = 1, \ldots, N. \tag{4.9}$$

The constant term $\tilde{\mathbf{C}} \in \mathbb{R}^N$ (4.7) lumps together all the terms that are linear in the velocity fluctuations $\boldsymbol{u}' \sim \boldsymbol{\phi}_i$ and contains the production/dissipation generated by the interaction between the fluctuating field and the mean flow. The linear term $\tilde{\mathbf{L}} \in \mathbb{R}^{N\times N}$ (4.8), containing all the terms quadratic in $\boldsymbol{u}'$, describes the production/dissipation generated by the interaction of the fluctuations with the mean flow and the viscous dissipation generated by the interactions of the fluctuation between themselves. Interestingly, the spectral properties of this term can be used as a proxy of the viscous dissipation present in the model. This property will play a key role for the temporal stabilisation of the system in the set up of the *a priori* sparsification technique. Lastly, the quadratic term $\tilde{\mathbf{Q}} \in \mathbb{R}^{N\times N\times N}$ (4.9) describes the nonlinear interactions of pairs of modes and it is the main responsible of the energy flow between modal structures. The study of the structure of this term is the main object of our analysis. Finally, the mass matrix $\mathbf{M} \in \mathbb{R}^{N\times N}$, with entries $M_{ij} = (\boldsymbol{\phi}_i, \boldsymbol{\phi}_j)$, takes into account the fact that the spatial modes may not be orthogonal and is introduced here for generality.

If the N modes span collectively an N-dimensional subspace $det(\mathbf{M}) \neq 0$. Thus $\mathbf{M}$ is invertible and the system (4.6) can be rearranged as

$$\dot{a}_i(t) = C_i + \sum_{j=1}^{N} L_{ij}a_j(t) + \sum_{j=1}^{N}\sum_{k=1}^{N} Q_{ijk}a_j(t)a_k(t) \qquad i = 1\ldots, N, \tag{4.10}$$

with

$$C_i = \sum_{q=1}^{N} M_{iq}^{-1}\tilde{C}_q, \qquad L_{ij} = \sum_{q=1}^{N} M_{iq}^{-1}\tilde{L}_{qj} \quad \text{and} \quad Q_{ijk} = \sum_{q=1}^{N} M_{iq}^{-1}\tilde{Q}_{qjk}. \tag{4.11}$$

If the set of modes $\boldsymbol{\phi_i}$ are orthonormal, as for the POD. Then, $M_{ij} = \delta_{ij}$ holds, and equation (4.10) is formally identical to equation (4.6). As observed by Rempfer and Fasel (1994a), the infinite dimensional matrix $\mathbf{M}$ should be first inverted and then truncated to maintain a good prediction accuracy. However, for all the models (POD and DFT based) considered in this thesis we have observed that the mass matrix $\mathbf{M}$ has a strong diagonal structure. Hence, the error performed by truncating it to size (N, N) and then inverting it can be reasonably assumed to be small.

Since the spatial modes satisfy automatically the boundary conditions, the expansion (4.3) provides a suitable foundation to examine interactions between coherent structures in complex geometries. Here, we follow established approaches (Rempfer and Fasel, 1994b)) and analyse such interactions by introducing the modal energies $e_i(t) =$

$\frac{1}{2}a_i(t)a_i(t)$, $i = 1, \ldots, N$. The instantaneous rate of change is given by

$$\dot{e}_i(t) = C_i a_i(t) + \sum_{j=1}^{N} L_{ij} a_i(t) a_j(t) + \sum_{j=1}^{N}\sum_{k=1}^{N} Q_{ijk} a_i(t) a_j(t) a_k(t), \qquad i = 1, \ldots, N, \tag{4.12}$$

obtained by multiplying (4.10) by $a_i(t)$. Note that, in a general case where the modes do not form an orthonormal set, the domain integral of the kinetic energy of velocity fluctuations is not the straightforward sum of the terms $e_i(t)$ and is given by

$$E(t) = \frac{1}{2}\int_{\Omega} \mathbf{u}'(t, \mathbf{x})^2 d\Omega = \frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{N} M_{ij} a_i(t) a_j(t). \tag{4.13}$$

The right hand side of equation (4.12) is composed of three terms describing energy transfers between the hierarchy of modes. The first two describe variations of energy due to production/dissipation arising from interactions with the mean flow and from viscous effects (Noack et al., 2011). The third term, involving the tensor $\mathbf{Q}$ describes the nonlinear energy flows in modal space regulated by the interactions of triads of modes $a_i(t)a_j(t)a_k(t)$. How to exploit the structure of this term for turbulent flow realisation to improve the performances of reduced order models is the main focus of this work. Additional insight on the long time behaviour of energy flows can be gained taking the temporal average of equation (4.12). The first term right hand side of the average of (4.12) is zero since the temporal coefficients $a_i(t)$ have zero mean. Similarly, at the left hand side, the average of the temporal derivative of the density of kinetic energy $\dot{e}(t)$ defined as

$$\overline{\dot{e}(t)} = \int_0^T \dot{e}(t)dt = \frac{1}{T}(e(T) - e(0)) \tag{4.14}$$

is zero if the energy of the system remains bounded and the length of the average interval T increases. The resulting time average energy budget of the system is regulated by

$$\sum_{j=1}^{N} L_{ij}\overline{a_i a_j} + \sum_{j=1}^{N}\sum_{k=1}^{N} Q_{ijk}\overline{a_i a_j a_k} = 0, \quad i = 1, \ldots, N. \tag{4.15}$$

If the temporal modes are uncorrelated in time $\overline{a_i a_j} \sim \delta_{ij}$ the first term left hand side simplifies in $L_{ii}\overline{a_i a_i}$. Equation (4.15) shows that in average the energy balance of the flow is regulated by the production dissipation described in the linear part and by the nonlinear energy transfer described by the quadratic part. As explained in Balajewicz et al. (2013) the average total power has to vanish for a model with $N \to \infty$. However, truncated Galerkin models are expected to predict an excess of kinetic energy making (4.15) positive. This excess of energy is classically absorbed by a eddy-vicosity model (Noack et al., 2008, 2005, 2011). More specifically, in the literature two different approaches have been proposed. First in the framework of the *a posteriori* stabilisation the idea proposed by Cordier et al. (2010) of automatically calibrating the coefficients of the system by means of linear regression. The second approach, proposed by Balajewicz et al. (2016), consists in rotating the basis into a more dissipative configuration modifying the eigenvalues of **L** as a mean to calibrate the correct viscous dissipation.

To systematically analyse triadic interactions we follow Rempfer and Fasel (1994a). Triadic interactions are defined in a time averaged sense by the triadic interaction tensor $\mathbf{N} \in \Re^{N \times N \times N}$ with entries defined as

$$N_{ijk} = Q_{ijk}\overline{a_i a_j a_k}, \tag{4.16}$$

where the overbar denotes temporal averaging. For later convenience we define the global amount of energy transferred to each mode by the triadic interactions as:

$$T_i = \sum_{j=1}^{N}\sum_{k=1}^{N} Q_{ijk}\overline{a_i a_j a_k}. \tag{4.17}$$

This term results from the Galerkin projection of the nonlinear term of the Navier-Stokes equations, $-\mathbf{u} \cdot \nabla \mathbf{u}$. It can be shown that,

$$\int_{\Omega} \mathbf{u} \cdot (\mathbf{u} \cdot \nabla \mathbf{u}) d\Omega = 0 \tag{4.18}$$

where the boundary terms resulting from integration by parts are equal to zero in case of periodic or confined flows. However, as pointed out by (Loiseau and Brunton, 2018), in several flow configurations this term can be assumed small and neglected for sake of simplicity. In modal space (4.18) translates in the following energy conservation property

$$\sum_{i=1}^{N} T_i = \sum_{i=1}^{N}\sum_{j=1}^{N}\sum_{k=1}^{N} Q_{ijk}\overline{a_i a_j a_k} = 0. \tag{4.19}$$

This physically means that, on average, no energy is produced or dissipated in the system during the cascade process from the large to the small scales of motion. The role of the nonlinear term in equation (4.12) is to transfer energy between modes describing different scales. As observed by (Loiseau and Brunton, 2018; Balajewicz et al., 2016) the conservation of this property throughout the sparsification procedure is of paramount importance for the physical meaningfulness of the sparse model.

The methodology introduced here is purely mathematical in nature, *i.e.* no physics is incorporated in the process. In addition, spatial modes obtained from classical decompositions have generally global support over the domain (see e.g. Taira et al. (2017)). The result is that the evolution equations (4.10) are not strictly sparse in the sense employed by Brunton et al. (2016). In fact, unless particular symmetries apply, the tensor **Q** is generally dense, *i.e.* most of its entries Q_{ijk} are different from zero and the right hand side of (4.10) contains all monomial terms in the modal amplitudes $a_i(t)$ up to order two. However, the key aspect motivating this work is the observation that in turbulent realisations of the Navier-Stokes equations only a subset of all triadic interactions contributes to a significant degree to the overall energy budget (Couplet et al., 2003; Rempfer and Fasel, 1994b). In this sense, sparsity is a primarily an *a posteriori* feature of solutions, *i.e.* a feature of the quadratic interaction tensor $\mathbf{N}$. The approach developed in this work uses this observation and aims to generate Galerkin models with a sparse coefficient tensor $\mathbf{Q}^s$. This sparse approximation of the original tensor $\mathbf{Q}$ should preserve information about the original dynamics of the system in the sense that the mismatch between the original $\mathbf{N}$ and the sparse transfer tensor $\mathbf{N}^s$ is as small as possible across the hierarchy of modes. This property ensure that although the coefficients tensor are sparse the relevant dynamics of the flows (contained in N) is retained throughout the sparsification procedure.

4.2 *A posteriori* Sparsification

The first sparsification methodology we discuss, involves an *a posteriori* evaluation of the relative strength of energy interaction between a predetermined set of flow structures defined by spatial modes. The sparsification occurs pruning from the right hand side of the system (4.6) the coefficients that do not contribute significantly to the dynamics while modifying the retained ones to keep the reconstruction error, computed with respect the modal acceleration, as low as possible. To our goal we use a procedure akin to that utilised in previous work for calibrating Galerkin models from data (Perret et al. (2006); Cordier et al. (2010); Xie et al. (2018)) and more recently for the identification of sparse dynamical systems from data(Brunton et al., 2016).

Starting from a velocity field $\mathbf{u}(t, \mathbf{x})$ a Galerkin model (4.6) is generated as shown in figure 4.1 panels (a) and (b). The main idea is to recast the Galerkin system (4.6), shown in panel (b), in a matrix from suitable to perform linear regression (as sketched in panel-(c)). Since N_t snapshots of the velocity field are available from simulation we can arrange samples of the temporal coefficients $a_i(t_j)$, $i = 1, \dots, N$ and $j = 1, \dots, N_t$, into the data matrix $\mathbf{A} \in \Re^{N_t \times N}$, with entries $A_{ij} = a_i(t_j)$. Each column of $\mathbf{A}$ contains the time history of a given temporal mode a_i for a time span equal to the simulation time. Similarly, we construct the modal acceleration matrix $\dot{A} \in \Re^{N_t \times N}$, each column containing the time derivative of a given temporal coefficients. Finally, exploiting the polynomial structure of the Galerkin system (4.10) its terms right hand side can be rearranged in the database matrix $\mathbf{\Theta}(\mathbf{A}) \in \Re^{N_t \times q}$

$$\mathbf{\Theta}(\mathbf{A}) = \begin{pmatrix} 1 & a_1^1 & a_2^1 & \dots & a_N^1 & a_1^1 a_1^1 & \dots & a_N^1 a_N^1 \\ \vdots & \vdots & \vdots & & \vdots & \vdots & & \vdots \\ 1 & a_1^{N_t} & a_2^{N_t} & \dots & a_N^{N_t} & a_1^{N_t} a_1^{N_t} & \dots & a_N^{N_t} a_N^{N_t} \end{pmatrix}, \tag{4.20}$$

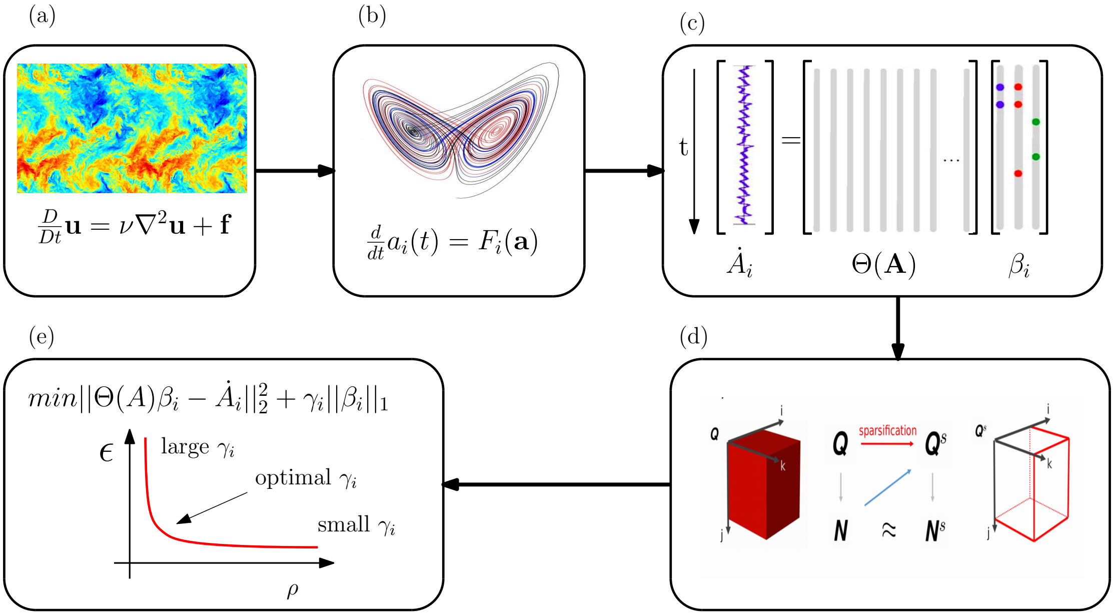

FIGURE 4.1: Schematic idea of the *a posteriori* sparsification procedure. Starting from a PDEs system in physical space (panel-(a)), this is transformed in a set of ODEs in modal space by Galerkin projection (panel-(b)). Once obtained the temporal evolution of the modes and their derivative the reduced order system (in panel (b)) can be recast in matrix form as shown in panel (c). Panel-(e) shows the idea of sparsification *i.e.* seeking for a sparse approximation of Q, Q^s such that $N^s \sim N$. Sparsification is obtained solving the LASSO problem for increasing values of the regularisation weight γ_i, obtaining a family of different models plotted on the $\rho - \epsilon$ plane as shown in panel-(e).

called nonlinear feature library in Brunton et al. (2016). The number of columns equal to $q = (N+1) + N(N+1)/2$ corresponds to number of features, *i.e.* the sum of constant, linear and quadratic interactions. Each column contains the temporal evolution of these interactions. The number of quadratic coefficients is only $N(N+1)/2$ instead of N^2 because the interaction between mode i and j is considered only once in (4.20). This avoids the columns of $\mathbf{\Theta}(\mathbf{A})$ to become linearly dependent, which would in turn result in numerical stability issues in the solution regression problem (see e.g. Perret et al. (2006) and Cordier et al. (2010)). As shown in figure 4.1-(c) arranging the projection coefficients tensors $\mathbf{C}$, $\mathbf{L}$ and $\mathbf{Q}$ associated to the i-th mode into a coefficient vector $\boldsymbol{\beta}_i \in \Re^q$, the Galerkin system (4.6) can be equivalently expressed as

$$\dot{\mathbf{A}}_i = \mathbf{\Theta}(\mathbf{A})\boldsymbol{\beta}_i, \quad i = 1, \dots N, \tag{4.21}$$

where $\dot{\mathbf{A}}_i$ is the $i-$th column of the modal acceleration matrix defined previously. An important aspect in the set up of (4.21) is the computation of the modal acceleration as two different approaches are possible. The first consists in computing the acceleration $\partial_t\mathbf{u}$ inside the CFD solution loop for every time step by backward finite differences. The acceleration is then written to disk for the prescribed time-steps. A second approach consists in computing $\partial_t\mathbf{u}$ as the right hand side of (4.2)-(a). This does not require any online computation inside the solution loop but it has been observed to produce higher reconstruction error with respect to the finite difference based approach. Thus, we chose to follow the first methodology.

The key idea is that if some nonlinear interactions are more important than others, then the corresponding entries of the coefficient vector $\boldsymbol{\beta}_i$, and consequently $\mathbf{Q}$, can be shrunk to zero with minor effects on the structure of the triadic interaction tensor $\mathbf{N}$ and consequently on the predictive ability of the resulting model. This is conceptually sketched in figure 4.1-(c) where the matrix form of a Galerkin system (4.21) for the Lorenz attractor is shown. Since only a small subset of interactions participates to the dynamics only a small subset of the entries of $\boldsymbol{\beta}_i$ is selected (coloured dots in panel (c)). This produces a sparse approximation of $\mathbf{Q}$ without affecting the properties of $\mathbf{N}$ as sketched in figure 4.1-(d).

The challenge is to find a systematic method to identify the dominant interactions and prune unnecessary coefficients whilst calibrating the remaining model coefficients such as to preserve the overall energy budget. In this work the established sparsity-promoting regression technique known as LASSO regression (Least Absolute Shrinkage Selection Operator, see Tibshirani (1996)) is employed. This algorithm can be applied to (4.21) column by column separately leading to a set of N optimisation problems of the form

$$\min_{\boldsymbol{\beta}_i} ||\boldsymbol{\Theta}(\mathbf{A})\boldsymbol{\beta}_i - \dot{\mathbf{A}}_i||_2^2 + \gamma_i||\boldsymbol{\beta}_i||_1, \qquad i = 1, ..., N, \tag{4.22}$$

one for each mode, where $\|\cdot\|_p$ denotes the l_p norm of a vector. The first least squares error term (called also l_2 part) in the objective function in (4.22) produces calibrated models that have minimum prediction error about the dynamics of the system, represented by the modal acceleration (see discussion in Cordier et al. (2010) and Couplet et al. (2005)). The second term (called l_1 part) penalises large model coefficients, regularises the regression and encourages sparsity in the solution by shrinking exactly to zero coefficients in $\boldsymbol{\beta}_i$, corresponding to columns of $\boldsymbol{\Theta}(\mathbf{A})$, with little dynamical influence. Ideally, to prune unnecessary coefficients, a penalisation term proportional to the cardinality of $\boldsymbol{\beta}_i$, $\text{card}(\boldsymbol{\beta}_i)$, would formally be more correct (Jovanović et al., 2014). However, the resulting optimisation problem would be computationally intractable even for Galerkin models of modest dimensions. In fact, this penalisation is usually relaxed to the computationally tractable l_1 term (Ramirez et al., 2013). Regardless, the optimisation problems (4.22) are convex and thus have an unique solution. In addition, the approach lends naturally to parallelisation, since the optimisation problems can be solved independently for each mode. In initial stages of the research, we have found approaches based on sequential thresholded least-squares (Brunton et al. (2016); Zhang and Schaeffer (2019); Loiseau and Brunton (2018)) to be not sufficiently robust. Hence, solutions of (4.22) have been computed using the `sklearn` (Pedregosa and Varoquaux (2011)) library, which implements a sub-gradient descent algorithm to manage the non differentiability of the l_1 norm.

The weights γ_i in equation (4.22) are arbitrary and can be tuned to trade prediction ability (when they are small) for sparsity (when they are large). To formalise these concepts, we introduce the global reconstruction error ϵ

$$\epsilon = \sum_{i=1}^{N} \frac{||\boldsymbol{\Theta}(\mathbf{A})\boldsymbol{\beta}_i - \dot{\mathbf{A}}_i||_2^2}{||\dot{\mathbf{A}}_i||_2^2} \tag{4.23}$$

and the system density ρ

$$\rho = \frac{1}{Nq} \sum_{i=1}^{N} \mathrm{card}(\boldsymbol{\beta}_i). \tag{4.24}$$

In equation (4.23), the absolute reconstruction error $||\boldsymbol{\Theta}(A)\boldsymbol{\beta}_i - \dot{\mathbf{A}}_i||_2^2$ is normalised with the mean squared acceleration $||\dot{\mathbf{A}}_i||_2^2$ to balance the global reconstruction error across the hierarchy, which would be otherwise dominated by the most energetic modes. On the other hand, the density ρ ranges from 0, when all interactions have been pruned, to 1, for a fully connected model. Note that for large models, the density is dominated by the quadratic tensor $\mathbf{Q}$. As shown in 4.1-(e) solving problem (4.22) for different values of the regularisation weights γ_i leads to a one-parameter family of models, producing a Pareto front (Schmidt and Lipson, 2009) on the ρ–ϵ plane. More specifically, we expect to obtain systems with large densities and small value of the reconstruction error for small values of γ_i and the opposite for large values of γ_i. Both these two scenarios have small practical utility since in the first case no sparsification is performed while in the latter the model is not able to provide accurate predictions. However, since only a subset of triadic interactions is relevant, the expectation is that a sweet spot appears on this curve, defining 'optimal' penalisation coefficients γ_i such that it is possible to reach low values of the density (*i.e.* several coefficients are eliminated) and, at the same time, keeping the reconstruction error low.

In addition, we observe that the penalisation coefficient γ_i can be chosen independently for each index i, implying that reconstruction error and sparsity can be modulated arbitrarily across the spectrum of modes. In our analysis we consider two different modulation strategies. In strategy S1, the weight is constant for all modes, $\gamma_i = \gamma$. This strategy sparsifies more aggressively the equations of motion of low-energy modes, because the l_1 penalisation term has a higher importance than the l_2 component. We also introduce strategy S2, where the weight is normalised with respect to the mean squared modal acceleration as $\gamma_i = ||\dot{\mathbf{A}}_i||_2^2\gamma$. This is equivalent to solving problem (4.22) using the relative error in (4.23) as least-squares component of the objective function. This strategy results in a more balanced sparsification across the hierarchy of modes and avoids earlier truncation, *i.e.* when all coefficients of a high-index mode are set to zero. Other strategies can be, of course, devised. Here, we mention, for instance, the possibility to tune the penalisation coefficients to obtain a uniform sparsification across the spectrum or to obtain a uniform relative reconstruction error. Analysing these strategies is an interesting avenue for future work.

One potential modification of this approach is discussed in Loiseau and Brunton (2018), namely to enforce that the nonlinear term in the sparsified Galerkin model conserves energy exactly (see e.g. Balajewicz et al. (2013) and Noack et al. (2008) for a formal definition). In practice, this can be achieved by introducing to the problem (4.22) a set of constraints between the coefficient vectors $\boldsymbol{\beta}_i$ belonging to different modes.

$$\boldsymbol{D}\mathbf{B} = \boldsymbol{d}. \tag{4.25}$$

Where the matrix $\mathbf{B} = vec([\boldsymbol{\beta}_1, \quad \ldots \quad, \boldsymbol{\beta}_N])$ is the vectorized form of the sparse matrices of the coefficients $\boldsymbol{\beta}_i$. The matrices $\mathbf{D}$ and $\mathbf{d}$, containing only 0 and 1, describe linear relations between pairs and triads of coefficients $\boldsymbol{\beta}_i$ belonging to different modes i. These matrices have a sparse structure and can be efficiently stored in memory with packages for the management of sparse matrices. More importantly, if linear energy constraints are added the whole problem made up by (4.22) and (4.25) remains convex and can be solved using iterative methods. The philosophy of these approach consist in considering large models made up of dozens of modes. This fact with the addition of these constraints, coupling together the regression problems of all modes, generates one optimisation problem of much larger dimension. This makes the solution of the coupled LASSO problem computationally expensive. As we will demonstrate later in the paper, the energy conservation error of models obtained from the unconstrained approach is small in relative terms. This occurs because the temporal coefficients in the matrix $\mathbf{A}$ are originally obtained from an energy conserving nonlinearity, and the regression 'discovers' this property from data. Hence, in this work we always solve problems (4.22) independently, without additional constraints. This confirms the importance of having a training dataset $\mathbf{A}$ and $\dot{\mathbf{A}}_i$ of good quality.

Finally, one last consideration about the regression set up. One of the key conceptual aspects of this work is that we aim to sparsify equations that are mathematically dense, *i.e.* the entries of the matrices in the system (4.6) are mainly nonzero, but whose realisations (solutions) are sparse, *i.e* the energy transfer occurs according preferential patterns. This brings the question whether to perform sparsification on the energy equation (4.12) or on the momentum equation (4.6). It can be shown (see appendix C) that the solution of the LASSO problem (4.22) with respect to the momentum equation or the energy equation leads to the same solution $\boldsymbol{\beta}_i$. Therefore, there is no need to reformulate problem (4.22) in the terms of the energy equation to sparsify energy interactions.

4.3 *A-priori* sparsification

The *a posteriori* sparsification involves a prune and tune procedure on the coefficients of (4.6) according to the energy relevance of the corresponding triadic interactions. In this approach the modal basis from which the Galerkin model is generated remains fixed during the sparsification while the numerical value of the triadic interactions is modified. Consequently, the link between modal structures and energy transfer could be easily lost in the process leading to a questionable interpretation of the physics of the modes. The *a priori* sparsification aims to fill this gap seeking a new modal basis that automatically generates, by Galerkin projection, models with a sparse triadic interaction tensor **Q**. Broadly speaking, we follow the idea of looking for this new basis as linear combination of the original one used to generate the model. In this work the POD basis will be used as starting point for its good energy reconstruction properties but other choices are possible.

4.3.1 Subspace Rotation Technique

To identify the new set of modal structures, we utilise the subspace rotation technique introduced in the context of stabilisation of Galerkin-based reduced order models by Balajewicz et al. (2013) (see also Amsallem and Farhat (2012)). Geometrically, this consists in seeking a rotation of an N-dimensional POD subspace within a larger POD subspace of dimension M. The rotation is defined by a transformation matrix $\mathbf{X} \in \Re^{M\times N}$, satisfying $\mathbf{X}^\top\mathbf{X} = \mathbf{I}$ to ensure that the rotated basis functions remain an orthonormal set. The rotated basis functions and the associated temporal coefficients, denoted in what follows with a tilde, are expressed as a linear combination of the original POD spatial and temporal modes via the rotation matrix X_{ji} as

$$\tilde{\boldsymbol{\phi}}_i(\mathbf{x}) = \sum_{j=1}^{M} X_{ji}\boldsymbol{\phi}_j(\mathbf{x}), \tag{4.26a}$$

$$\tilde{a}_i(t) = \sum_{j=1}^{M} X_{ji}a_j(t). \tag{4.26b}$$

It is worth pointing out that finding a new set of modal structures directly would be a much higher dimensional optimisation problem to tackle, with size proportional to the numbers of modes required times the number of degrees of freedom of the problem at hand (equal to the grid point times the dimension of the problem). Seeking new modal structures as a linear combination of POD modes represents a significant reduction in complexity, which can be finely controlled by varying the dimension M. In addition, using POD modes, under the assumption that a small rotation is sufficient to promote sparsity , has the advantage of producing a basis inheriting a good energy reconstruction properties.

The coefficients of the Galerkin system (4.6) projected in the new subspace are then defined as

$$\tilde{Q}_{ijk} = \sum_{i,j,k=1}^{N} \sum_{p,q,r=1}^{M} Q_{pqr} X_{pi} X_{qj} X_{rk} \tag{4.27a}$$

$$\tilde{\mathbf{L}} = \mathbf{X}^\top \mathbf{L} \mathbf{X} \tag{4.27b}$$

$$\tilde{\mathbf{C}} = \mathbf{X}^\top \mathbf{C}. \tag{4.27c}$$

In equation (4.27), $\mathbf{C} \in \Re^M$, $\mathbf{L} \in \Re^{M\times M}$ and $\mathbf{Q} \in \Re^{M\times M\times M}$ are the Galerkin coefficient tensors obtained from the M-dimensional set of original POD modes. The goal is to find a particular rotation for which the rotated quadratic interaction coefficient tensor $\tilde{\mathbf{Q}}$, obtained from (4.9) with the rotated spatial modes $\tilde{\boldsymbol{\phi}}_i$, has a sparse structure, i.e. where many of quadratic interaction coefficients are identically zero. In reason of the practical intractability of the zero norm in optimisation (Jovanović et al., 2014), here we use the l_1 norm operator, denoted as $\|\cdot\|_1$, to measure the sparsity of the quadratic interaction coefficient tensor (Rubini et al., 2020b; Brunton et al., 2016). However, it is worth noting that any rotation is necessarily accompanied by a loss of fluctuation kinetic energy reconstructed by the new basis, since the original N POD modes are optimal in this respect by design. Therefore, it won't be possible to sparsify the system without affecting the ability of the new basis of reconstructing kinetic energy but a sparsity-energy reconstruction trade off will always be present. To quantify the energy loss associated with the rotation $\mathbf{X}$ the fluctuation kinetic energy reconstructed by a P-dimensional set of POD can be calculated by utilising the modal energies $\lambda_i = \overline{a_i a_i}$, $i = 1, \dots, P$. Arranging them into a diagonal matrix $\boldsymbol{\Lambda}_P \in \Re^{P\times P}$, the trace $Tr(\boldsymbol{\Lambda}_P)$ defines an upper bound for the reconstructed energy for any P-dimensional set of basis functions, due to well-known optimality property of POD. Similarly, the energy reconstructed by the rotated basis can be expressed the modal energies $\tilde{\lambda}_i = \overline{\tilde{a}_i \tilde{a}_i}$, $i = 1, \dots, N$, of the rotated temporal coefficients (4.27b) and arranging them into the diagonal matrix $\tilde{\boldsymbol{\Lambda}}_N = \mathbf{X}^\top \boldsymbol{\Lambda}_M \mathbf{X} \in \Re^{N\times N}$, where the expression at the right hand side follows directly from (4.27b). The loss of reconstructed fluctuation kinetic energy with respect to an N-dimensional POD subspace is then quantified as

$$\mathcal{J}(\mathbf{X}) = Tr(\boldsymbol{\Lambda}_N - \mathbf{X}^\top \boldsymbol{\Lambda}_M \mathbf{X}). \tag{4.28}$$

As observed in Balajewicz et al. (2013), the quantity $\mathcal{J}(\mathbf{X})$ is necessarily non-negative due to the optimality of the original POD basis, i.e.

$$\int_\Omega \|\mathbf{u}'(t, \mathbf{x})\|^2 \mathrm{d}\Omega \geq Tr(\boldsymbol{\Lambda}_N) \geq Tr(\tilde{\boldsymbol{\Lambda}}_N), \tag{4.29}$$

where the last equality holds if, e.g., $\mathbf{X} = \mathbf{I}_{M\times N}$ resulting in no rotation away the POD subspace. In addition, the quantity (4.28) is also identically zero when $M = N$, for any $\mathbf{X}$, since any linear combination of N POD modes necessarily spans the same original

N-dimensional subspace. Therefore, in this work only configurations with $M > N$ are considered.

Mathematically, the trade-off between energy optimality and sparsity is expressed by formulating the following constrained optimisation problem

$$\begin{aligned} \min_{X} \quad & Tr(\mathbf{\Lambda}_N - \mathbf{X}^\top \mathbf{\Lambda}_M \mathbf{X}) && \text{(4.30a)} \\ \text{subject to} \quad & ||\tilde{\mathbf{Q}}||_1 \leq \xi^{-1} ||\mathbf{Q}||_1, && \text{(4.30b)} \\ & \mathbf{X}^\top \mathbf{X} = \mathbf{I}, && \text{(4.30c)} \end{aligned}$$

referred to as problem P1 in what follows. The role of the objective function (4.30a) is to favour small transformation matrices that minimises the loss of energy optimality, producing a set of basis function with good representation ability, as in Balajewicz et al. (2016). On the other hand, the constraint (4.30b) encourages sparse solutions, with ξ an arbitrary penalisation parameter that controls the relative l_1 norm of the rotated quadratic coefficient tensor. Increasing the penalisation weight ξ can thus promote sparsity, because some entries of $\tilde{\mathbf{Q}}$ are shrunk to zero during the solution of (4.30) by the non-differentiability of the l_1 norm Tibshirani (2013); Friedman et al. (2008). As we will discuss later on, it is worth noting that problem (4.30) might not have any feasible solution if the penalisation weight is too large, as it might not be possible to find a transformation matrix that shrinks the l_1 norm of the rotated coefficient tensor by the required amount. This is consequence of the fact that the l_1 penalisation is applied not on the optimisation variables X_{ji} but on the entries of the rotated quadratic interaction tensor $\tilde{Q}_{ijk}$. An important characteristic of optimisation problem (4.30) is that while the objective (4.30a) is convex, the sparsity-promoting constraint (4.30b) is not, as it involves cubic polynomials in the optimisation variables, the entries of the rotation matrix X_{ij}. Consequently, the solution might no be unique and several local minima, corresponding to different sets of basis functions, may be obtained by starting the optimisation from different initial guesses. These aspects are problem dependent and they will be discussed later on when the test cases are introduced.

Assuming a feasible solution of problem (4.30) can be found, the Galerkin model constructed from projection onto the optimal rotated basis does not necessarily possess better long-term temporal stability characteristics than the original POD model. In fact, it is well-known that POD-Galerkin models exhibit long-term instability because of the deficit of energy dissipation attributed to the truncation of small dissipative scales (Noack et al., 2008; Schlegel and Noack, 2015). In the present case, the transformation $\mathbf{X}$ obtained from solution of (4.30) does not necessarily result in an improved description of dissipative processes. Classically, this issue is cured by introducing, *a posteriori*, an eddy-viscosity-type term in the Galerkin model (Galletti et al., 2004; Noack et al., 2005; Östh et al., 2014). However, an *a posteriori* correction would not remain in the spirit

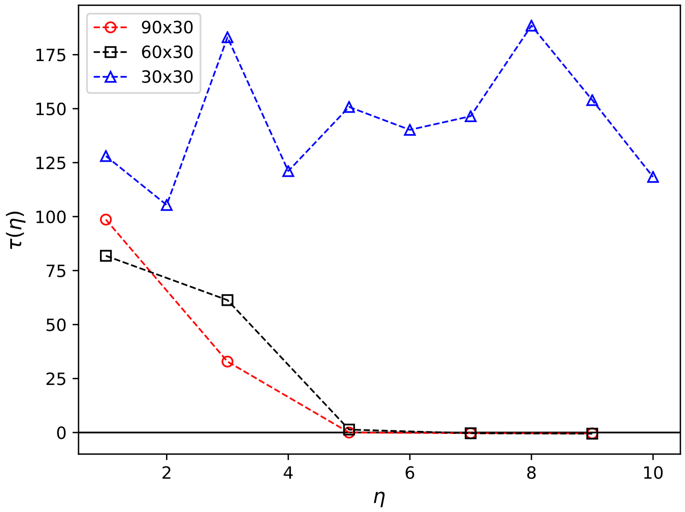

FIGURE 4.2: Excess of turbulent kinetic energy $\tau(\eta)$ against η.

of the present approach. We thus favour the subspace-rotation-based stabilisation approach proposed by Balajewicz et al. (2013, 2016), which can be introduced naturally in the present formulation. In practice, we augment problem (4.30) with the additional implicit constraint

$$Tr(\tilde{\mathbf{L}}) = Tr(\mathbf{X}^\top \mathbf{L}\mathbf{X}) = -\eta \tag{4.31}$$

where the auxiliary variable $\eta \in \Re^+$ is chosen such that

$$\tau(\eta) = \frac{\overline{E(t)} - \overline{E_{\mathrm{DNS}}(t)}}{\overline{E_{\mathrm{DNS}}(t)}} = 0, \tag{4.32}$$

i.e. that the relative difference of the average fluctuation kinetic energies from DNS and from numerical simulation of the new model vanishes. $\overline{E(t)}$ is computed as the time average of the turbulent kinetic energy $e_i(t)$ after time integration of the model obtained with the new basis functions. Since it is not generally true for quadratic models this quantity will be independent from the initial condition. In this work we chose as initial condition for the time integration the first time step of the statistically converged DNS solution. The variable η controls dissipation mechanisms in the Galerkin model by altering the eigenvalues of $\tilde{\mathbf{L}}$ and ensures long-term stability. Figure 4.2 shows the effects of η on the excess of kinetic energy (5.12) for three different ratio of M/N. The flow considered is the lid-driven cavity at Reynolds 2×10^4. It can be observed how increasing η the excess of kinetic energy converges towards the level obtained by DNS. More interestingly figure 4.2 shows that a ration $M/N > 1$ is necessary for the stabilisation of the system. This is physically explainable with the fact that smaller scales (with modal index $M > N$) to introduce the correct amount of dissipation in the system. As observed by Balajewicz et al. (2013), η is not known a priori, but can be found in an inner optimisation loop to ensure that the excess of turbulent kinetic energy defined by

$\tau(\eta)$ is zero. With this additional constraint, problem (4.30) becomes

$$
\begin{aligned}
\min_{X} \quad & Tr(\mathbf{\Lambda}_N - \mathbf{X}^\top \mathbf{\Lambda}_M \mathbf{X}) && \text{(4.33a)} \\
\text{subject to} \quad & ||\tilde{\mathbf{Q}}||_1 \leq \xi^{-1} ||\mathbf{Q}||_1, && \text{(4.33b)} \\
& Tr(\mathbf{X}^\top \mathbf{L} \mathbf{X}) = -\eta \quad \text{with} \quad \tau(\eta) = 0, && \text{(4.33c)} \\
& \mathbf{X}^\top \mathbf{X} = \mathbf{I} && \text{(4.33d)}
\end{aligned}
$$

In this formulation, denoted as P2 henceforth, there is still only one free parameter, ξ, which is used to trade sparsity for energy optimality. The solution algorithm, explained in Algorithm 1 table, can be divided in two main steps. First a reduced problem without the sparsity constraint (4.35b) to find η such that the excess of kinetic energy (5.12) is smaller than a certain threshold e_{TOL}. Once this condition on $\tau(\eta_m)$ is satisfied the stabilisation parameter $\eta = \eta_m$ is fixed and the complete problem (4.35) can be solved for a given value of ξ returning as output the rotation matrix X. If a solution of (5.12) is not found within the maximum number of iterations N_{max} the while loop stops.

To numerically solve the problem (4.35) the procedures implemented in the open-source package for nonlinear and non-convex optimisation `NLopt` (Johnson, 2014) were used. We utilised a solver implementing the Method of Moving Asymptotes (MMA) algorithm (Svanberg, 2014). This requires the gradient of the objective function and of the constraints. The key element to make the procedure viable is to evaluate the sparsity-promoting constraint and its gradient as efficiently as possible. An efficient algorithm to compute these two quantities with costs scaling as $\mathcal{O}(M^3N) + \mathcal{O}(M^2N^2) + \mathcal{O}(MN^3)$ at most (as opposed to a naive implementation which requires $\mathcal{O}(M^3N^3)$ operations for the computation of $\tilde{\mathbf{Q}}$ and its l_1 norm and $\mathcal{O}(M^4N^4)$ for the evaluation of its gradient with respect to the rotation $\mathbf{X}$) can be derived and its discussion is deferred later on in this chapter.

Once (4.35) is solved, a new set of temporal and spatial modes can be found according to (4.26) while a new set of coefficients for the reduced order model can be found from (4.27). Since the value of ξ is arbitrary, it can be varied to obtain models with different sparsity. To formalise this concept we introduce the model density, defined as

$$\rho = \frac{||\tilde{\mathbf{Q}}||_0}{||\mathbf{Q}||_0} \tag{4.36}$$

where the l_0 norm $||\mathbf{Q}||_0$ counts the nonzero elements of the tensor $\mathbf{Q}$. This can also be expressed as the average of the modal densities ρ_i as:

$$\rho = \frac{1}{N} \sum_{i=1}^{N} \rho_i \quad \text{with} \quad \rho_i = \frac{||\tilde{\mathbf{Q}}_i||_0}{||\mathbf{Q}_i||_0}. \tag{4.37}$$

Algorithm 1 Algorithm for the solution of problem (4.35).

```
Input  : Endpoints values $[\eta_a \eta_b]$ and $[\chi_0 \chi_n]$. Convergence tolerance $e_{TOL}$ on (5.12),
         maximum number of iterations $N_{max}$, ROM starting and arrival space dimensions ($M$ and $N$)
         and Galerkin system matrices associated with the first $M$ most energetic POD modes
         $\mathbf{Q} \in \Re^{M\times M\times M}$, $\mathbf{L} \in \Re^{M\times M}$ and $\mathbf{C} \in \Re^{M}$.
Output: Rotation matrix $\mathbf{X} \in \Re^{M\times N}$
while $N_{iter} < N_{max}$ do
    Fix $\xi$
    Fix a value of $\eta_m = (\eta_a + \eta_b)/2$
    Solve system
        $\min_X \quad Tr(\mathbf{\Lambda}_N - \mathbf{X}^\top \mathbf{\Lambda}_M \mathbf{X})$    (4.34a)
        $\text{subject to} \quad Tr(\mathbf{X}^\top \mathbf{L} \mathbf{X}) = -\eta_m,$    (4.34b)
        $\mathbf{X}^\top \mathbf{X} = \mathbf{I}$    (4.34c)
    (4.35) without (4.35b)
    Reconstruct system matrices
    Integrate numerically the Galerkin system
    Calculate the excess of kinetic energy $\tau(\eta_m)$
    if $\tau(\eta_m) < \epsilon$ then
        fix $\eta = \eta_m$
        Solve Complete problem
            $\min_X \quad Tr(\mathbf{\Lambda}_N - \mathbf{X}^\top \mathbf{\Lambda}_M \mathbf{X})$    (4.35a)
            $\text{subject to} \quad ||\tilde{\mathbf{Q}}||_1 \leq \xi^{-1} ||\mathbf{Q}||_1,$    (4.35b)
            $Tr(\mathbf{X}^\top \mathbf{L} \mathbf{X}) = -\eta \quad \text{with} \quad \tau(\eta) = 0,$    (4.35c)
            $\mathbf{X}^\top \mathbf{X} = \mathbf{I}$    (4.35d)
        Terminate algorithm and return X
    if $sign(\tau(\eta_c) = sign(\tau(\eta_a))$ then
        Set $\eta_a = \eta_c$
    else
        Set $\eta_a = \eta_b$
```

To express the loss of energy captured by the rotated basis, we also introduce the global energy reconstruction factor

$$e_N = Tr(\tilde{\mathbf{\Lambda}}_N)/Tr(\mathbf{\Lambda}_\infty) \tag{4.38}$$

where $\tilde{\mathbf{\Lambda}}$ and $\mathbf{\Lambda}$ are the matrices having on the diagonal the modal energies of the rotated and original basis, respectively. Since, the energetically optimal POD is used as starting point the global energy reconstruction factor is always strictly lower than 1 by definition.

The density and reconstruction factor are affected by the dimensions M and N and by the penalisation weight ξ. To characterise the effects of these parameters, we can display each system obtained with a different ξ onto the $\rho - e_N$ plane. For each value of M and N, a family of reduced order models with different density/energy reconstruction trade off can be generated increasing the values of the penalisation parameter in the constraint (4.35b). The geometrical (size M and N of the subspaces) and the physical properties of the problem are expected to affect the shape and the slope of these curves.

This visualisation in conceptually similar to the approach used in the a posteriori sparsification technique with two expected differences. First, it is not possible to obtain a flat part decreasing the density since every rotation away from the energetically optimal POD automatically results into a loss in the ability of the decomposition in reconstructing turbulent kinetic energy. Second, in the case of a posteriori sparsification we observed that it is possible to reach zero values of density increasing the regularisation weight γ_i to arbitrarily large values. We expect this not to be possible with the subspace rotation approach since, given M and N, it is not possible to arbitrarly decrease the value of the density of the triadic interaction tensor $\tilde{\mathbf{Q}}$ via a rotation. In addition, we expect that this will result into sparse systems with generally higher densities with respect to the ones obtained with the a posteriori approach.

4.3.2 Computational cost of the algorithm

The sparsity promoting constraint involves polynomial functions of the optimisation variables. An analytical expression of the gradient of this constraint with respect to the rotation $\mathbf{X}$ can be also obtained, enabling fast gradient-based optimisation to be utilised. In addition, a careful examination of the tensor operations involved in the computation of the constraint and its gradient shows that a significant reduction of the scaling of costs can be obtained by reorganising some computations, avoiding repeated operations (see Pfeifer et al. (2014)). For instance, a naive calculation of the quantity

$$\tilde{Q}_{ijk} = \sum_{i,j,k=1}^{N} \sum_{p,q,r=1}^{M} Q_{pqr} X_{pi} X_{qj} X_{rk} \tag{4.39}$$

required in the evaluation of the sparsity-promoting constraint (4.30b) takes $\mathcal{O}(M^3N^3)$ operations because of the presence of the six (one for each index) nested for loops. However, a careful evaluation of operation (4.39) shows that it can be expressed as the sum of three different step each involving only four nested loops. This is done by introducing the auxiliary variables $A_{qri} \in \mathbb{R}^{M\times M\times N}$ and $B_{rij} \in \mathbb{R}^{M\times N\times N}$ defined as follows:

$$A_{qri} = \sum_{i=1}^{N} \sum_{p,q,r=1}^{M} Q_{pqr} X_{pi} \tag{4.40a}$$

$$B_{rij} = \sum_{i,j=1}^{N} \sum_{p,q=1}^{M} A_{qri} X_{qj} \tag{4.40b}$$

$$\tilde{Q}_{rij} = \sum_{i,j,k=1}^{N} \sum_{r=1}^{M} B_{rij} X_{rk} \tag{4.40c}$$

This operations only takes $\mathcal{O}(M^3N) + \mathcal{O}(M^2N^2) + \mathcal{O}(MN^3)$ operations. This results can be easily inferred from the indexes involved in the evaluations of expression (4.40)(a),(b) and (c), respectively. The gradient of the sparsity promoting constraint with respect to the entries X_{mn} of the rotation matrix $\mathbf{X}$ requires the computation of

$$\frac{\partial ||\tilde{\mathbf{Q}}||_1}{\partial X_{mn}} = \sum_{i,j,k=1}^{N} \frac{\partial \tilde{Q}_{ijk}}{\partial X_{mn}} \mathrm{sign}(\tilde{Q}_{ijk}). \tag{4.41}$$

The gradients $\partial \tilde{Q}_{ijk} / \partial X_{mn}$ can be expanded using the Kronecker's delta δ_{ij} as

$$\frac{\partial \tilde{Q}_{ijk}}{\partial X_{mn}} = \sum_{p,q,r=1}^{M} \mathrm{Q}_{pqr}(X_{pi}X_{qj}\delta_{kn}\delta_{rm} + X_{pi}X_{rk}\delta_{jn}\delta_{qm} + X_{rk}X_{qj}\delta_{in}\delta_{pm}). \tag{4.42}$$

Overall, a naive computation of $\partial||\tilde{\mathbf{Q}}||_1/\partial\mathbf{X}$ would require $\mathcal{O}(M^4N^4)$ operation, since the tensor expressions involves a total of eight indices. A similar reorganisation of the operations leads to tensor operations over four tensor indices only, with computational costs similar to the one of the rotation of $\mathbf{Q}$.

4.4 Summary

This chapter explained the set up of the mathematical tools used to generate Galerkin models of turbulent flows, study energy flows between modal structures and perform model sparsification. To start with, the generation of reduced order models of turbulent flows is explained and subsequently, the two different sparsification methods are presented. The first method, called *a posteriori* sparsification method, belongs to the class of supervised learning techniques. It consists in recasting the Galerkin model into a matrix form suitable to perform linear regression. The key idea is leveraging the properties of the l_1 penalisation on the coefficient vector. The non differentiable nature of this norm allows extracting from large datasets the subset of relevant features. At the same time, its convex nature allows recasting the interactions selection problem into a convex optimisation problem that can be solved efficiently. A linear regression approach is of simple implementation and the results are easily interpretable, both these qualities make it a good algorithm to incorporate the existing physical laws in. However, this methodology has also two main drawbacks, one concerning the sparsification process and the other the physical interpretability of the results. First, since the value of the regularisation weight γ_i is unknown, the sparsification consists in progressively increasing its value obtaining a whole family of models with different sparsity-accuracy trade offs. Among all these possibilities the user must choose the one that best suits his/her needs. However, a well known drawback of the supervised learning methodologies is the impossibility to know the performance of these models beforehand without evaluating them singularly by temporal integration. Second, the *a posteriori* sparsification technique consists in pruning the least relevant energetic interactions while tuning the numerical value of the retained ones to minimise the reconstruction error. Consequently, the numerical value of the triadic interactions is modified while the spatial basis functions from which the model is generated are not. This aspect could lead to some interpretability issues about the magnitude of the energy transfer and the flow structures associated with them.

The *a priori* sparsification technique aims to fill this gap seeking a new modal basis such that the resulting model obtained by Galerkin projection posses already a sparse structure. This is accomplished by rotating the original POD basis in a larger subspace. The rotation matrix is obtained as the solution of an optimisation problem aiming to minimise at the same time the energy loss and the l_1 norm of the rotated quadratic interaction tensor $\tilde{\mathbf{Q}}$. Since both the triadic interactions and the basis functions are modified this approach resolves the physical inconsistency of the *a posteriori* one. Moreover, since the penalisation is not applied directly on the optimisation variable but on the rotated quadratic interaction tensor we expect it not to be possible to reach arbitrarily low values of the density but a minimal value of the density exists for each model. This could remove from the problem the choice of the regularisation weight, that can be chosen as the one corresponding to the lowest density such that the solution of the optimisation problem is acceptable. These advantages with respect to the *a priori* technique come at the cost of a higher computational cost and loss of uniqueness for the optimisation problem. However, as it will be shown later in the results section for initial guesses sufficiently near the identity matrix the solution tends to fall into the same minimum.

Chapter 5

Results: *a posteriori* sparsification

In this chapter the *a posteriori* sparsification technique is applied to different models of increasing physical complexity. First, the one dimensional PDE describing the Kuramoto-Sivashinsky system is introduced as a paradigmatic model of multi-scale, chaotic flows evolving in a periodic domain. This can be conveniently described in Fourier space. For this model, we focus on the ability of the LASSO of capturing the relevant interactions between flow scales and on the influence of the definition of the regularisation weight on the selection of energetic interactions. Second, the two-dimensional chaotic flow evolving inside a lid-driven square cavity is considered. This configuration is chosen since the wall-bounded domain and the simple Cartesian geometry allow a straightforward generation of reduced-order models of increasing physical complexity. The goal of this analysis is to quantify the capabilities of the sparse regression of preserving the physics under different flow conditions, such as increasing Reynolds numbers, and different modal representations. Lastly, a more practical application is considered. Namely, the *a posteriori* sparsification technique is used as a mean to calibrate reduced order models of separated flows around a NACA 0012 airfoil. This test case will be also used as a testbed for the library developed to generate reduced order models from numerical data defined on arbitrary computational grids. With this last analysis, we intend to evaluate the scalability of the sparsification methodology to real-life engineering problems.

5.1 One Dimensional PDE: Kuramoto-Sivashinsky equation

To start with, the Kuramoto-Sivashinsky model, as a paradigmatic chaotic system characterised by multi-scale dynamics, is examined. Due to the periodic boundary conditions, scales interaction can be conveniently studied in Fourier space. This system is used to understand the applicability of the sparse regression framework to high dimensional reduced-order models.

5.1.1 Mathematical model

The Kuramoto-Sivashinsky equation describes the evolution of a zero-mean scalar field $u = u(t, x)$ on a periodic domain of length L. The mathematical model is given by

$$\begin{aligned} \partial_t u &= -uu_x - u_{xx} - u_{xxxx} \\ u(t,x) &= u(t,x+L) \\ u(0,x) &= u_0(x), \end{aligned}$$

where the subscript indicates partial differentiation with respect to the variables x and t. The dynamics of chaotic solutions of this model is qualitatively similar to the one of a turbulent flow, where energy is injected at the large scales by the laplacian term with the inverted sign. Subsequently, energy is transferred towards smaller scales by the nonlinearity and finally dissipated by the hyper-viscosity, represented by the fourth derivative term.

Due to the periodic nature of the boundary conditions we can conveniently solve this model in Fourier space. Introducing $\alpha = 2\pi/L$, we decompose the solution field as

$$u(t,x) = \sum_{k=-N}^{N} u_k(t) e^{i\alpha x} \tag{5.1}$$

where $u_k = u^*_{-k}$ holds ($*$ indicates the complex conjugates), since $u(t, x)$ is real. In Fourier space we obtain the following set of ODEs

$$\dot{u}_k = (\alpha^2 k^2 - \alpha^4 k^4) u_k - \frac{1}{2} i\alpha k \sum_{p=-N}^{N} u_p u_{k-p} \qquad k = 1, ..., N \tag{5.2}$$

The set of ODEs (5.2) is formally equivalent to the system (4.6) with the matrix coefficients computed according the expression (4.7), (4.8) and (4.9). Due to the symmetry of the domain and the Fourier expansion (5.1) of the velocity field, these matrices have a compact structure in modal space. More specifically, C_i is identically zero due to the zero mean field, $L_{ij} = (\alpha^2 k^2 - \alpha^4 k^4)\delta_{ij}$ and $Q_{kpq} = i\frac{k}{2}$ if $p + q = k$ and zero otherwise. Interestingly, we observe that the linear term L_{ij} is diagonal with entries defined by $(\alpha^2 k^2 - \alpha^4 k^4)$ and the amplitude of the mode itself u_k. In addition, it can be observed that $(\alpha^2 k^2 - \alpha^4 k^4) > 0$ if $0 < k < \alpha^{-1/2}$ and smaller than zero otherwise. This generates a turbulence-like energy cascade where energy is injected at low wave-numbers (large scales) and dissipated at the higher wave-numbers (small scales). The energy is transferred by nonlinear interactions between triads of modes satisfying the aforementioned relation between wave-numbers.

The system (5.2) has been numerically integrated using a fourth order Runge-Kutta time marching technique, where for the evaluation of the linear term pseudo-spectral utilities FFT() are appropriate. Figure 5.1-(a) shows $u(t, x)$ of a statistically converged

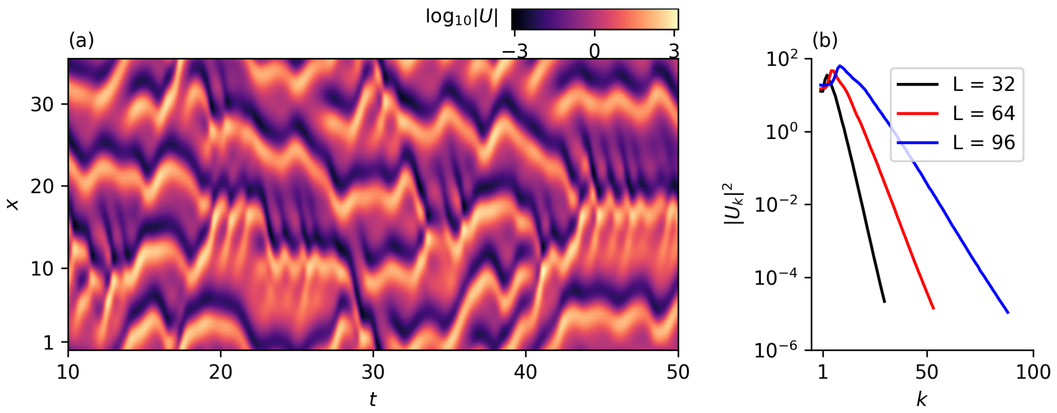

FIGURE 5.1: Temporal evolution of a solution of the Kuramoto-Sivashinsky equation obtained for $L = 32$ in panel-(a). Panel-(b): power spectrum of three different solutions for $L = 32$,64 and 96, respectively.

solution obtained for a domain length of $L = 32$, where we have set $N = L = 32$ to obtain simulations with a dynamical range spanning several orders of magnitude. It can be observed the formation of random wave-like structures repeating themselves inside the domain and evolving in time. In addition, we observe that after a long enough time interval all information about the initial condition is lost making its choice irrelevant for the present analysis. Panel (b) shows the power spectrum of three solutions obtained with different lengths of the domain L and with a number of modes such that $N = L$ in non-dimensional units. It can be observed that the energy contained at each scale initially grows, reaching its maximum corresponding to $k = \frac{1}{\sqrt{2}\alpha}$ (*i.e.* the maximum of the energy production function $\alpha^2 k^2 - \alpha^4 k^4$ and then decreases exponentially after that. More importantly, the range of dynamically relevant scales widens when the width of the domain L is increased. This happens because an elementary flow structure can repeat itself more when as the domain widens. Therefore, the parameter L is equivalent to the Reynolds number, *i.e.* increasing the length of the domain increases the range of dynamically active scales in the solution. This makes the model a good test-bed to evaluate the performances of the sparse regression on systems where the scale separation in the system widens as the complexity grows. In addition, to compare results obtained with different systems we use as a resolution $N = L$, to resolve similar dynamical ranges, for all the three different configurations considered.

To provide a stronger foundation to the sparsification results the sparsity pattern obtained with the LASSO regression is compared with the pattern of energy interaction inside the model. The kinetic energy of the k-th mode in the Fourier space is defined as

$$E_k(t) = \frac{1}{2} u_k(t) u_k^*(t), \tag{5.3}$$

and its temporal evolution is described by

$$\dot{E}_k = T_k + (\alpha^2 k^2 - \alpha^4 k^4) E_k. \tag{5.4}$$

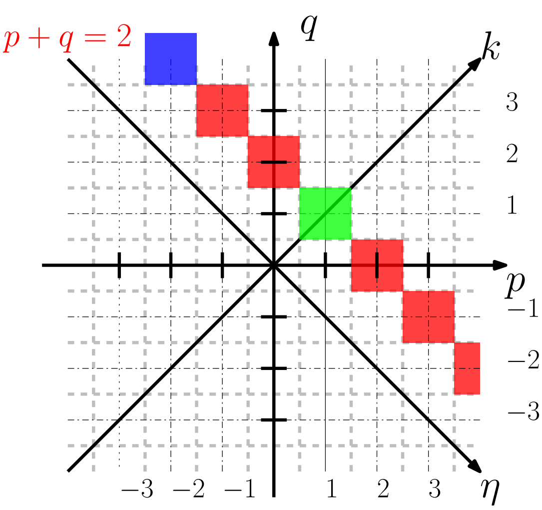

FIGURE 5.2: Visualisation of triadic interactions for wavenumber $k = 2$ on the $p - q$ or $k - \eta$ plane. In addition, green and blue squares show local and distant interactions, respectively.

The second term at the right hand side is a production/dissipation linear term. The term T_k is defined as

$$T_k(t) = Im\left\{\sum_{p+q=k} r_{pqk}\right\} = Im\left\{\sum_{p}(\alpha k)u_p u_{k-p} u_k^*\right\}, \tag{5.5}$$

where Im indicates the imaginary part of a complex number, and represents the sum of all triadic interactions r_{pqk} between wavenumbers k, p, q such that $k = p + q$. Because of the periodic boundary conditions T_k satisfies the energy preserving property

$$\sum_{k=1}^{N} T_k = 0. \tag{5.6}$$

5.1.2 Energy Analysis

When defined in Fourier space, the Kuramoto-Sivashinky equation (5.2) has a sparse structure. The linear term is diagonal while the nonlinear interactions only include modes such that the relation $p + q = k$ is satisfied. To aid the understanding, in the following sections, we visualise all nonlinear interactions (living inside the third order tensor Q_{pqr}) on a two dimensional plane, as illustrated in figure 5.2. Each point on this plane, defined by the coordinates (p, q) or (k, η), corresponds to a triadic interaction r_{pqk}. The axes (k, η) are defined as $k = p + q$ and $\eta = p - q$ respectively. We distinguish local and non-local interactions by their distance from the axis k. Points near the axis k corresponding to $p \sim q \sim k$ and small absolute value of η, *i.e.* the green square is an example of local interaction, while points such that $p >> q \sim k$ or $p \sim q >> k$, far from the k axis (large absolute value of η), such as the blue square, are classified as

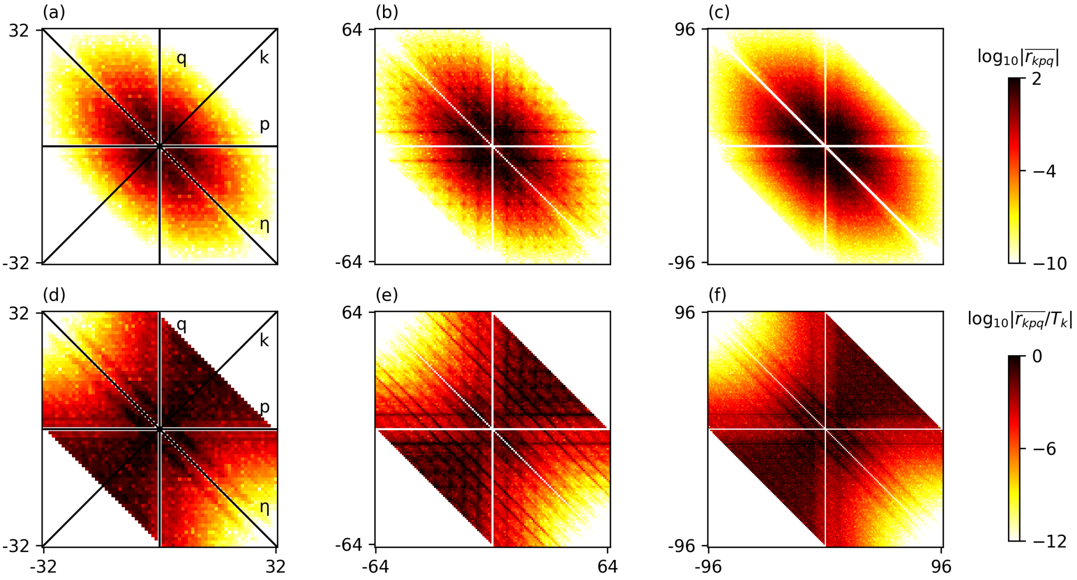

FIGURE 5.3: Average absolute value of energy interactions $\overline{r}_{pqk}$ for systems with $L = 32, 64$ and 96 in panels (a),(b) and (c), respectively. Average absolute value of the relative energy interactions $\overline{r}_{pqk}/T_k$ for systems with $L = 32, 64$ and 96 in panels (d),(e) and (f), respectively

nonlocal interactions. In what follows, large-large scales interactions are identified as points near the origin of this diagram, and large-small scales interactions as points near the bottom-right and upper-left corners.

Panels (a, b, c) of figure 5.3 show the base ten logarithm of the average absolute energy transfer term $|\overline{r}_{pqk}|$ for three different systems with $L = 32,64$ and 96, respectively. It can be observed that the most energetic interactions are large-large scale interactions located near the origin of the axes. The large-small scales and the small-small scales interactions have less relevance. This structure of the interactions suggests that a good sparsification strategy could be based on the elimination of the interactions with high frequency modes from the low frequency ones and on the complete elimination from the system of the high frequency low energy modes. Panels (d,e,f) of figure 5.3 show the average of the absolute energy transfer normalised by the total energy transfer associated with that mode, the quantity $\overline{r}_{pqk}/|\overline{T_k}|$. This plot illustrates the relative importance of the interactions within each mode. It can be observed that for the small scales (large k, away from the main diagonal), all interactions are equally important in relative terms. This means that energy interactions are non-local for this system since the dynamics of the small scales depends on all the larger scales and not just those that are immediately adjacent in Fourier space.

5.1.3 Effect of the regression on the coefficients

The sparse structure of the Kuramoto-Sivashinsky model in Fourier space allows to easily analyse the trend of the coefficients as the sparsification algorithm (4.22) is applied on the system (5.2). In this section, the effects of the l_1 penalisation on the model coefficients are analysed. We consider the behaviour of the coefficients of the nonlinear term for three different Fourier modes located in different parts of the spectrum having wavenumber k equal to 1,10 and 20, respectively. These coefficients are reconstructed by solving problem (4.22) for increasing values of the regularisation weight. The numerical value of the coefficients are shown in panel (a,b,c) of figure 5.4. Due to the simple structure of this model all the interactions contributing to the dynamics of a certain wavenumber k live on the pseudo-diagonals of diagram shown in figure 5.2. This structure facilitates the visualisation of Q_{pqr} as a function of η. Each panel of figure 5.4 shows the effect of the sparsification procedure on the Fourier modes taken into consideration. The red dashed line corresponds to the analytic value of the coefficients equal to $\frac{\alpha k}{2}$ while the black lines are the reconstructed coefficients obtained with four different increasing values of the regularisation weight γ_k. Interestingly, it can be observed how for the modes with low wavenumber ($k = 1$ and $k = 10$ in panels (a) and (b), respectively), describing large and more energetic structures, increasing the regularisation weight shrinks to zero the coefficients corresponding to distant interactions (high absolute value of η). This is in agreement with the shape of the absolute energetic interactions r_{kpq} displayed in panels (a,b,c) of figure 5.3, meaning that the dynamics of the modes with low wavenumbers is not affected by the interactions with modes with high wavenumber (distant interactions). Conversely, panel-(c) shows the coefficients of the mode with wavenumber $k = 20$, located in the mid-rear part of the spectrum

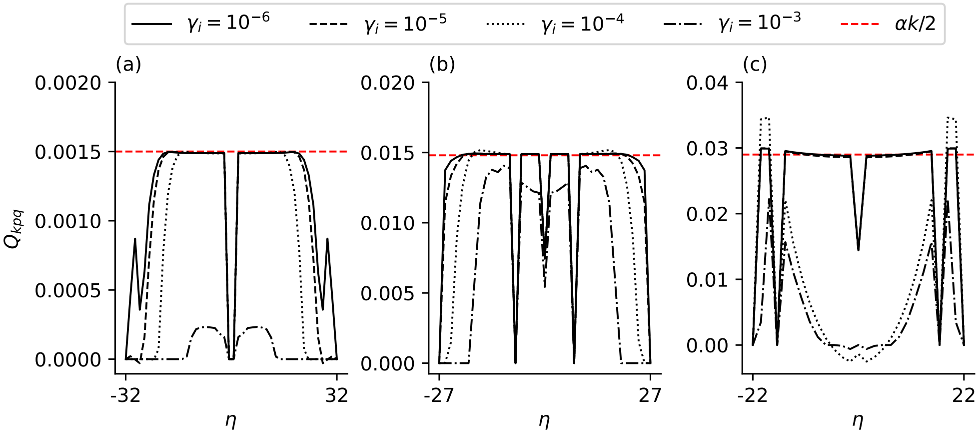

FIGURE 5.4: Reconstructed Q_{pqr} with $L = 32$ for $k = 1$, $k = 10$ and $k = 20$ in panels (a),(b) and (c), respectively. The red line corresponds to the analytical value, $\alpha k/2$, while the black lines correspond to models obtained with increasing regularisation weights.

for the model considered. It can be observed how in this case the coefficients eliminated first are the ones with a low value of η, corresponding to interactions occurring between modes of similar size, *i.e.* local interactions. More specifically, as the regularisation weight is increased only the interactions between low and high wavenumbers (*i.e.* between large and small structures) are retained. This interactions are located near the p and q axis in the diagram 5.2 and correspond to distant or non-local interactions. This result is in agreement with the shape of the relative energy interactions shown in figure 5.3 panels (d) to (f) where the energy transfer in the high-energy modes are dominated by the nonlocal interactions. More interestingly, it can be observed that the curves in panel (b), even in the central part, are not perfectly flat, as expected from the analytical solution, but present a M-like shape where the values of the coefficients first decreases in the central part because of the effect of the l_1 regularisation and then slightly increases before dropping to zero. This behaviour is due to the optimiser trying to balance the absence of some coefficients through a slight modification of the remaining ones to keep the reconstruction error flat at least for low regularisation weights. Arguably, this balancing mechanism is one of the key aspects that makes the l_1 regression-based approach superior in terms of accuracy to a simple truncation or an "a priori physics informed" approach where the interactions are manually eliminated with modifying the retained ones. The effect of the modification of the retained coefficients on the stability and on the energy conservation properties of the sparse system is an aspect of paramount importance when more complex configurations are considered later.

5.1.4 Model sparsification

The sparsification procedure consists in solving repeatedly problem (4.22) for increasing values of the regularisation weight γ_k to obtain systems characterised by different sparsity patterns. The set-up of the regression problem is totally analogous to what shown in chapter 4 with the only difference of a modification of the database matrix according to the mode considered due to the property of the Fourier modes $u_k = u^*_{-k}$ that can create collinearity between the columns of $\boldsymbol{\Theta}$ if all the modes are considered ad once. To evaluate the prediction accuracy of the sparsified system, we define a reconstruction error based on the difference between the original rate of change $\dot{u}_k$ and the one obtained by the sparsified system $\dot{\tilde{u}}_k$ (denoted by a tilde). The quadratic reconstruction error can be easily defined for Fourier modes as:

$$\epsilon = \frac{\overline{\sum_{k=1}^{N} (\dot{\tilde{u}}_k - \dot{u}_k)^2}}{\overline{\sum_{k=1}^{N} \dot{u}_k^2}}. \tag{5.7}$$

To evaluate the sparsification performances of the LASSO applied to this system the reconstruction error ϵ is displayed as a function of the density of the system ρ. The density can be defined as $\rho = N_{!zero}/N_{tot}$ where $N_{!zero}$ is the number of interactions different from zero and N_{tot} is the total number of interactions originally different from zero. These quantities are mathematically defined as the l_0 norm of the coefficients vector $\boldsymbol{\beta}_k$ restricted on the active set of coefficients of $\mathbf{L}$ and $\mathbf{Q}$ according to (4.24). We can also define the density relative to each mode $\rho_k = N^i_{!zero}/N^i_{tot}$. Both these quantities range from $\rho = 0$ (fully sparse system) to $\rho = 1$ (a fully dense system). To understand how different definitions of the regularisation weight γ_k (strategy S1 or S2) affect the selection of the interactions, it is of interest to monitor the behaviour of ρ_k as the global density ρ decreases.

5.1.5 l_1-based regression with constant regularisation weight

Figure 5.5 is constructed by varying the regularisation weight uniformly across all modes (strategy S1). For each value of the regularisation weight, which is defined as $\lambda_i = \tilde{\lambda}$ we first solve the optimisation problem and then evaluate the reconstruction error ϵ and density ρ associated with the sparsified model. To better visualise the effect of the increase of the regularisation weight on the triadic interactions we define the tensor $\boldsymbol{\epsilon}$ whose entries ϵ_{ijk} are coloured with the value of the reconstruction error ϵ, as defined in (4.23), obtained when the corresponding interaction coefficient Q_{pqr} is eliminated from the system. Due to the compact structure of the Kuramoto-Sivashinsky equation in Fourier space, we can visualise the tensor on the $p - q$ plane as explained in figure 5.2. The result is visualised in panels of figure 5.5 showing the tensor $\boldsymbol{\epsilon}$ and the corresponding interactions patterns for systems with $L = 32$, 64 and $L = 96$, respectively. Interestingly, regardless of the size of the problem two common trends emerge.

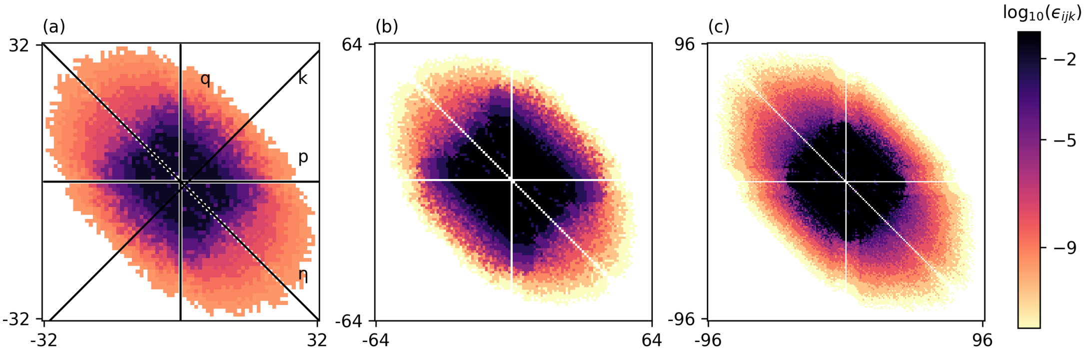

FIGURE 5.5: Sparsity pattern displayed by the contour of the reconstruction error ϵ obtained with a sparsification approach with regularisation weight $\lambda_i = \tilde{\lambda}$ constant across i. Panels (a), (b) and (c) show $L = 32$, 64 and 96, respectively.

First, the most energetic modes are sparsified by removing first the distant interactions, points located in the upper left and lower right corners of the graph. Therefore, the LASSO sparsification produces, for the larger structures of the flow, a LES-like truncation. This means that an LES approach is optimal in the sense of problem (4.22). Second, high frequency and low energy modes are completely eliminated from the system as the regularisation weight increases producing truncation. Interestingly, increasing the size of the domain L we observe that, at least for the high energy modes, the number of retained interactions is almost constant moving towards high values of k. This suggests that, at least for the low wave-number high energy modes, the interactions are local in modal space, with the bandwidth of the active band dependent on the physics of the problem.

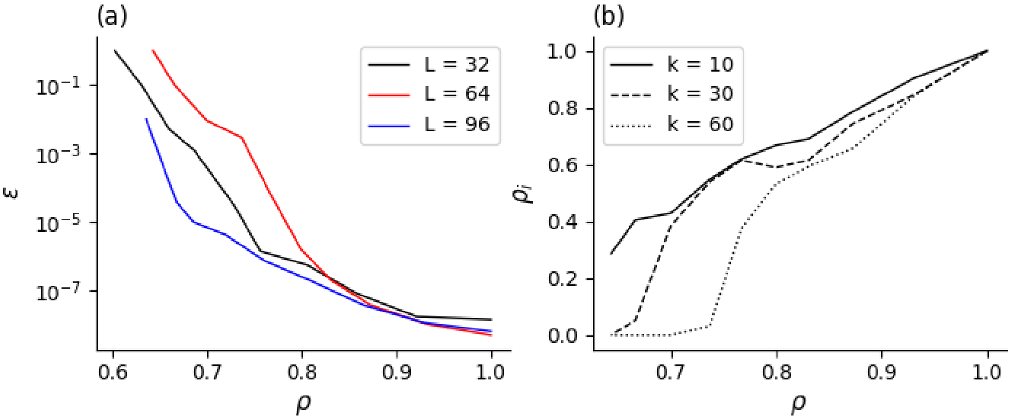

FIGURE 5.6: Panel-(a): reconstruction error ϵ vs density ρ for the constant regularisation weight approach. Panel-(b): plot of the modal density ρ_i against the global density ρ for different modes for the model $L = 64$.

Problem (4.22) contains the free parameter γ_i. Thus, for the same configuration, a family of different systems can be generated by varying the value of the regularisation weight. To each value of γ_i corresponds two different values of ρ and ϵ. The corresponding model can be displayed on the $\epsilon - \rho$ plane. Figure 5.6-(a) shows the $\epsilon - \rho$ curves for the three systems with different domain lengths previously considered. A plateau in the right part of the graph can be observed, with a knee point located around $\rho = 0.85$. This suggests that it is possible to remove interactions from the system affecting only marginally its predictive accuracy. This occurs because other interaction coefficients are modified in the optimisation to preserve the overall energy balance. Interestingly, we observe that the curves are not much affected by the size of the domain. This can be also qualitatively observed in figure 5.5, showing that the sparsification patterns have a self-similar structure on the $k - \eta$ plane. We argue that this is due to the self-similar shape of the spectrum (5.1-(b)) and to the constraint $N = L$ leading to the same energy resolution for all the systems.

As discussed in the methodology section the regularisation weight γ_k can be also varied mode by mode leading to different sparsification patterns in modal space. When the sparsification weight is kept constant the pattern of the selected interactions follows the one of the absolute energy interactions leading to the sparsification of the most energetic modes and to the truncation of the least energetic ones. This behaviour can be better visualised by plotting the modal density ρ_k against the global density ρ as shown in figure 5.6-(b). Here, three modes in different regions of the spectrum for the system with $L = 64$ are shown. As expected from a visual analysis of the identified sparsity pattern shown in figure 5.5 the high energy mode ($k = 10$) is sparsified without eliminating all the interactions since ρ_k never reaches zero. On the other hand, when considering low energy/high wavenumber modes ($k = 30, 60$) ρ_k drops to zero when the global density reaches $\rho \sim 0.75$. This means that the evolution of these modes is eliminated from the system. In addition, as a consequence of this definition of the regularisation weight we observe that generally the high index modes can be sparsified more efficiently than the low index ones. Therefore, this definition of the regularisation weight results in a trade off between sparsification of the most energetic modes and truncation of the least energetic ones. This also provides the best result in terms of the reconstruction error since in the formulation of problem (4.22) the l_1 penalisation term interacts minimally with the least-squares part responsible for the reconstruction accuracy of the reconstructed system.

However, the main philosophy of a sparsification procedure is to generate a cheaper model without performing any truncation but only eliminating the least important interactions within the dynamics of each mode.

5.1.6 l_1-based regression with variable regularisation weight across modes

To avoid truncating the high-wavenumber, low-energy, modes we vary the value of the penalisation weight across the spectrum we scale the value of γ_k in equation (4.22) with the square root of the average kinetic energy of each mode, *i.e.* as $\gamma_k = \sqrt{\overline{u_k u_k^*}}\tilde{\lambda}$, where $\tilde{\lambda}$ is an arbitrary constant. We have chosen this approach with the idea of introducing a penalisation that is proportional to the total average energy transfer. This formulation is equivalent to solve problem (4.22) with a least-squares component equal to the relative reconstruction error as defined in (4.23). However, the choice of the scaling of the regularisation weight is not unique and other choices can be explored.

Figure 5.7 shows the tensor ϵ obtained with this definition of the regularisation weight. Similarly to figure 5.5 each cell, corresponding to a triad of modes p, q, k is coloured according to the value of the reconstruction error obtained eliminating that set of interactions from the system. Interestingly, the sparsity pattern identified by the LASSO with this definition of the regularisation weight resembles the pattern of relative energy transfers shown in panels (d,e,f) of figure 5.3. With this definition of the regularisation weight we observe that for the low index high energy modes the distant interactions are eliminated first. This results is similar to what observed for constant regularisation weight with the only difference that these modes are sparsified much more aggressively. It will be shown later how this affects the slopes of the sparsification curves. A different behaviour can be observed for the low energy high wavenumber modes

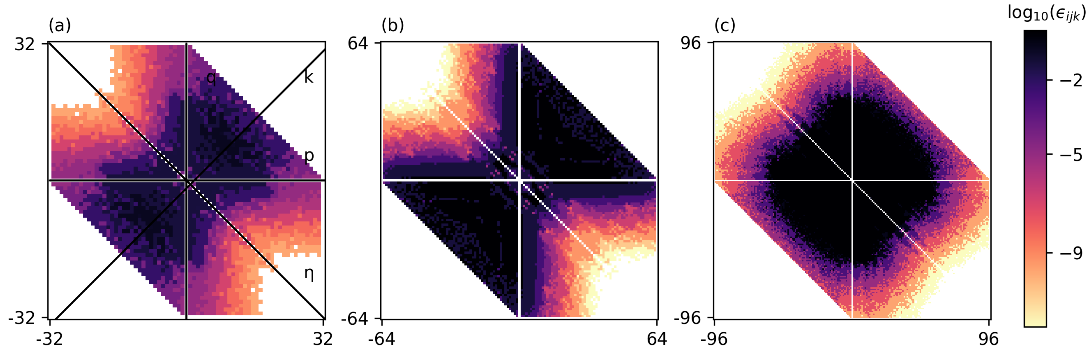

FIGURE 5.7: Sparsity pattern displayed by the contour of the reconstruction error ϵ obtained with a sparsification approach with regularisation weight defined as $\lambda_k = \sqrt{u_k u_k^*}\tilde{\lambda}$ constant across i. Panels (a),(b) and (c) correspond $L = 32$,64 and 96, respectively.

whose interactions are almost completely retained with only the elimination of some local interactions (located near the k axis). The overall result is a system that is properly sparsified without truncating any mode originally present in the dense system.

Figure 5.8-(a) shows the reconstruction error as a function of the density. It can be observed that the plateau is small or non existent and the error grows more rapidly than in the case where the penalisation weight is kept constant across the modes. This is due to the fact that, for a given target density, it is more convenient to truncate the system, *i.e.* remove completely the small scales, than sparsifying more aggressively large scale/high energy modes. To analyse the sparsification pattern across modes, figure

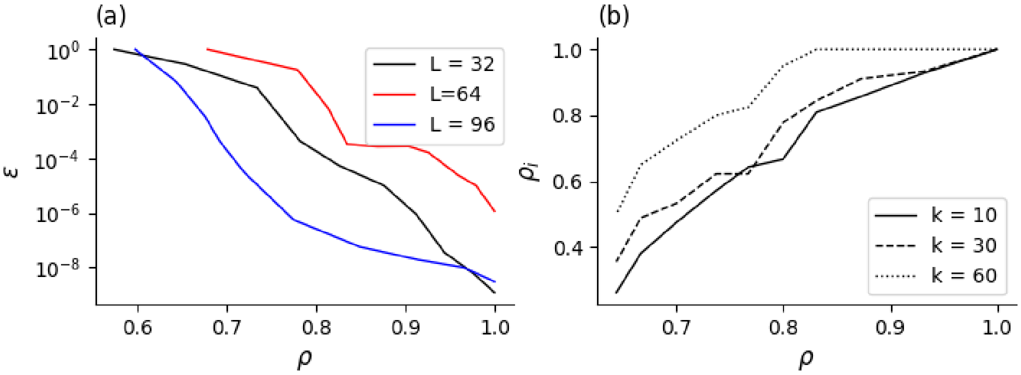

FIGURE 5.8: Panel-(a): reconstruction error ϵ vs density ρ for the variable regularisation weight approach. Panel-(b): plot of the modal density ρ_i against the global density ρ for different modes for the model with $L = 64$.

5.8-(b) shows the modal density ρ_k against the global density ρ for the system with $L = 64$. As qualitatively observed in figure 5.7-(b) no mode is completely eliminated from the system during the sparsification. Interestingly, conversely to what observed in figure 5.8-(b) with this definition of γ_k the large scales/high energy modes are more effectively sparsified. This is in agreement with the shape of relative energetic interactions shown in 5.3-(e) and it is consequence that, in relative terms, only the local interactions are important for the low indexes mode while all the interactions are equally important for the high indexes modes. Finally, we observe that generally the approach with constant regularisation weight (strategy S1) produces generally lower reconstruction error and sparsification curves with a more favourable L-shape with respect to the approach S2.

5.2 Lid-Driven Cavity flow

We now move on to the analysis of two-dimensional unsteady chaotic flow inside a lid-driven square cavity. This is an established test case for the development and validation of model order reduction techniques (Cazemier et al., 1998; Terragni et al., 2011; Balajewicz et al., 2013; Arbabi and Mezić, 2017; Fick et al., 2018), and we thus consider it here as an example to demonstrate the ideas discussed in the introduction and to test the mathematical framework on a more geometrically complex test case with respect to the Kuramoto-Sivashinsky model. The lid-driven cavity flow is particularly suitable for our purposes since the simple domain, contoured by no-slip walls, allows a straightforward computations of the coefficients matrices of the reduced order-model (4.6). At the same time, by increasing the Reynolds number, models of arbitrary complexity can be easily obtained. Results shown in this section have been published in (Rubini et al., 2020b,a).

5.2.1 Flow field

The Reynolds number is defined as $Re = LU/\nu$, where L and U are, respectively, the cavity edge and the lid velocity while the kinematic viscosity is denoted by ν. The edge length and the lid velocity are set equal to 1 in non-dimensional units. Therefore, different values of the Reynolds number can be easily obtained just by varying the kinematic viscosity. The domain is defined in nondimensional Cartesian coordinates $\mathbf{x} = (x, y)$ and the velocity field is defined by the components $\mathbf{u} = (u, v)$. In this first analysis, we are mainly interested in understanding the capabilities of the l_1 sparsification framework to preserve the physical properties of the flow throughout the sparsification procedure. Thus a flow configuration with $Re = 2 \times 10^4$ where dynamics is chaotic, as shown by Auteri et al. (2002), was chosen. The simulations are performed in `OpenFOAM` using a modified version of the solver `icoFOAM` that also outputs the snapshots of the Eulerian acceleration $\partial_t \mathbf{u}(t, \mathbf{x})$ to compute the modal acceleration $\dot{a}_i$ needed in (4.22). The convective and viscous terms are spatially discretised with a second order finite volume technique and the temporal term with a semi-implicit Crank-Nicholson scheme. Special treatments of the singularities developing at the top corners due to the

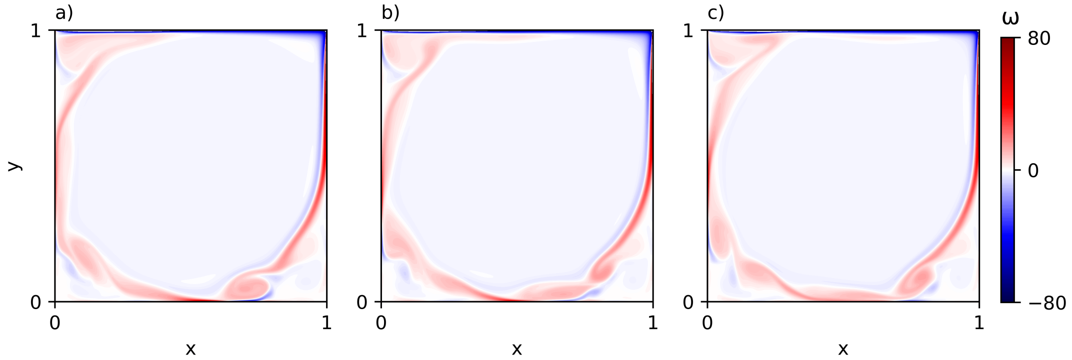

FIGURE 5.9: Vorticity field ω of three different snapshots separated by one non-dimensional time unit, increasing from left to right.

discontinuity in the velocity boundary conditions (Botella and Peyret, 1998) were not deemed necessary. Three snapshots of the vorticity field obtained from these simulations are shown in figure 5.9. Most of the dynamically interesting features in this regime originate at the bottom-right corner of the cavity. Specifically, the secondary vortex in the recirculation zone is shed erratically, producing wave-like disturbances advected along the shear layer bounding the primary vortex. The characteristic non-dimensional frequency of this wave-like motion is $f = 0.7$. From simulation, we extract $N_t = 1500$ velocity snapshots using a nondimensional sampling period $\Delta t = 0.1$. These settings are sufficient to adequately time-resolve the fast scales as well as to include many shedding events at the bottom right corner, making the regression problems (4.22) statistically reliable.

5.2.2 Modal Decomposition

To examine the role of the subspace used for projection on the sparsity of triadic interactions in modal space, we consider and compare two families of Galerkin models obtained by projecting the Navier-Stokes equations onto the subspaces generated by two different modal decompositions. First, we consider models generated using POD modes. It is well known that POD produces compact reduced order models, but has the shortcoming of mixing together fluid motions at different temporal/spatial scales but having similar energy content (Mendez et al., 2019). Second, we consider models generated from modes oscillating at a single frequency obtained from a procedure that is equivalent to a Discrete Fourier Transform (DFT) of the velocity snapshots. For practical convenience, we obtain the two distinct sets of modes using the same computational technique, based on the approach proposed by Sieber et al. (2016) which only operates on the temporal correlation matrix. Briefly, the method considers the temporal correlation matrix $\mathbf{R} \in \mathbb{R}^{N_t \times N_t}$, with entries

$$R_{ij} = \frac{1}{N_t} \left(\mathbf{u}'(t_i, \mathbf{x}), \mathbf{u}'(t_j, \mathbf{x}) \right), \tag{5.8}$$

and then defines a filtered correlation matrix $\mathbf{S}$, with entries

$$S_{ij} = \sum_{k=-N_f}^{k=N_f} g_k \mathbf{R}_{i+k,j+k} \tag{5.9}$$

given by the application of the filter coefficient vector $\mathbf{g}$ along the diagonals of the correlation matrix. A set of temporal coefficients, ordered according their energy content, $\mathbf{a}_i = [a_i(t_1), \dots, a_i(t_{N_t})]$ and associated mode energies λ_i is then obtained from the eigendecomposition of $\mathbf{S}$,

$$\mathbf{S}\mathbf{a}_i = \lambda_i \mathbf{a}_i, \tag{5.10}$$

so that $\mathbf{a}_i^\top \cdot \mathbf{a}_j = \lambda_i \delta_{ij}$. As discussed in Sieber et al. (2016), when the filter is extended over the entire dataset and in the limit of number of samples tending to infinity, the filtered correlation matrix converges to a Toeplitz, circulant matrix. Then, its eigenvalues trace the power spectral density of the underlying data set. On the other hand, the eigenvectors $\mathbf{a}_i$ correspond to the Fourier basis. This procedure generates conjugate pairs of modal structures with same energy oscillating at a single frequency. These can be viewed as a set of modal oscillators exhibiting periodic fluctuations (Taira et al., 2017) and tracing fluid motion at on a two-dimensional subspace. In practice, for a finite-length dataset, we filter the temporal correlation matrix assuming periodicity using a box filter, as suggested in Sieber et al. (2016). Hereafter, we will refer to the modal structures identified by this procedure as DFT modes.

One important consideration is that, unlike Dynamic Mode Decomposition (see Rowley et al. (2009); Schmid (2010); Chen et al. (2012)), DFT lacks the ability to discern and identify dominant frequency components. Instead, a number of modes equal to the number of snapshots utilised is produced, oscillating in conjugate pairs at specific frequencies determined by the sampling period Δt and observation time T (Mendez et al., 2019). This property, *picket fencing*, results in frequencies that are integer multiples of the fundamental frequency $f_1 = T^{-1}$, up to the Nyquist component $f_{Nyq} = (2\Delta t)^{-1}$. In addition, unlike for POD, as the length of the dataset is increased, the number of energy-relevant modes increases and low-frequency modes with little dynamical importance appear. The approach we use here is to divide the dataset into five partition of thirty time units, covering an average of 20 cycles of the dominant oscillatory component, and providing sufficient frequency resolution to distinguish small scale spectral features. In addition, two possible ways of sorting pairs of modal structures are possible, *i.e.* by energy content (using the eigenvalues λ_i) or by frequency. Models obtained with the two sorting schemes will be referred to as DFT$_e$ and DFT$_f$, respectively.

n	1	5	10	15	20	26	50	75	80	95	100	300
POD	0.26	0.74	0.85	0.89	0.9	0.95	0.98	0.99	0.995	0.998	0.999	-
DFT$_e$	0.17	0.49	0.62	0.74	0.8	0.88		0.92	0.95	0.97	0.98	1
DFT$_f$	0.02	0.07	0.15	0.39	0.43	0.47	0.91		0.96	0.97	0.98	1

TABLE 5.1: Normalised cumulative energy distribution $e(n)$ for POD and DFT modes, where the latter are sorted by energy content (DFT$_e$) or by frequency (DFT$_f$).

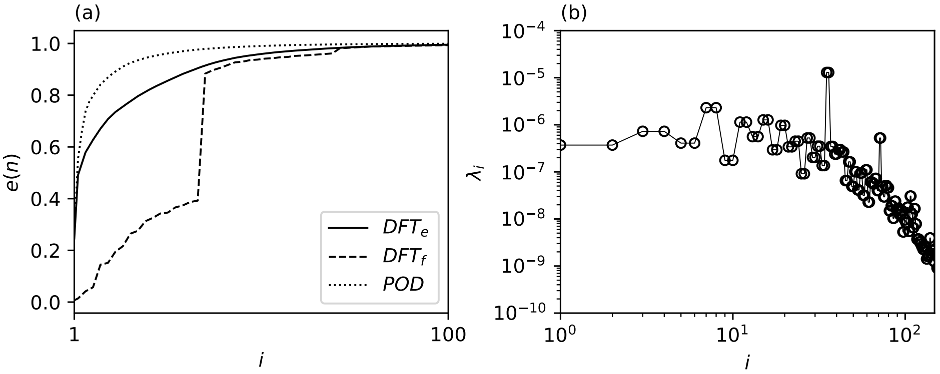

FIGURE 5.10: Panel (a): cumulative sum of the first 100 eigenvalues of S_{ij} for the three decompositions considered. Panel (b): distribution of the modal energies of DFT modes sorted by frequency.

We now focus on the characteristics of the modal structures obtained by these two methods. We denote the normalised cumulative sum of the eigenvalues λ_i of the (filtered) correlation matrix as

$$e(n) = \sum_{i=1}^{n} \lambda_i / \sum_{i=1}^{N_t} \lambda_i, \tag{5.11}$$

describing the fraction of the fluctuation kinetic energy captured by the first n elements of the expansion (4.3). This quantity is shown in figure 5.10-(a) for the POD and the two possible DFT sorting schemes. As expected, a larger energy is captured by the POD basis. For the DFT decomposition, the energy-based sorting is more efficient at data compression, although the difference vanishes for large n, since for low energy POD modes converge towards Fourier modes and the two sorting schemes are equivalent. The modal energies associated to the DFT$_f$ modes are shown in figure 5.10-(b) as a function of the modal index i. The distribution is characterised by a continuous component, with modal energy decaying with frequency, and a discrete component, with a fundamental peak for the pair of modes (31, 32) and its first few harmonics. The peak, at a non-dimensional frequency $f = 0.7$, is physically originated from the high-energy structures transported along the shear layer by the rotation of the main vortical structure.

This can be observed in panels (a) and (b) of figure 5.11, showing the vorticity field ω of the DFT mode pair (31, 32). This pair of modes describes a vorticity perturbation

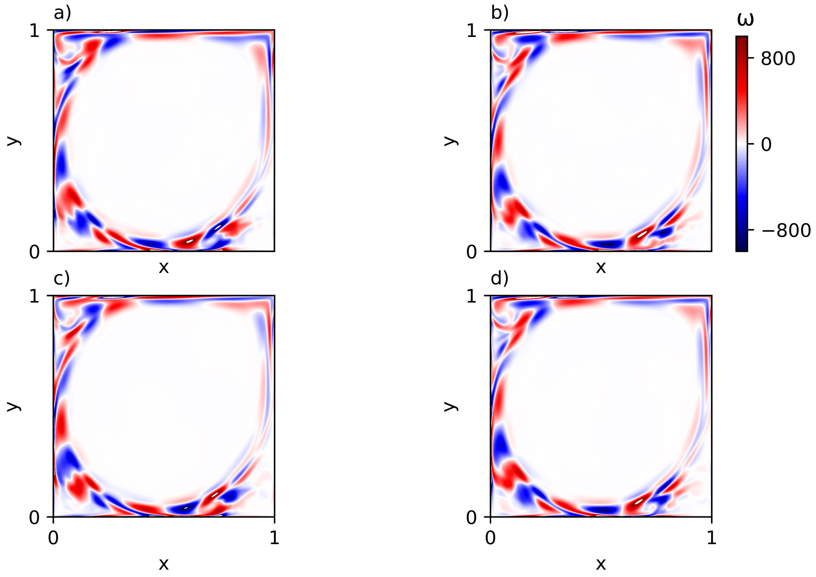

FIGURE 5.11: Vorticity field of the most energetic pair of DFT modes, panels (a) and (b), and of the first two POD modes, panels (c) and (d).

having the form of a wave travelling along the edge of the main vortex. Hence, the spatial structure of the two modes is shifted in the direction of the shear layer by half wave. Travelling-wave structures in cavity flows have already been observed in simulation by Poliashenko and Aidun (1995); Auteri et al. (2002) and characterised by global stability analysis and Koopman analysis by Boppana and Gajjar (2010) and Arbabi and Mezić (2017), respectively. The two leading POD modes, reported in panels (c) and (d) of figure 5.11, have the same energy and capture the same travelling-wave pattern described by the leading DFT mode pair.

5.2.3 Energy Analysis

To provide a more robust foundation for the understanding of the sparsification results, we first focus on the analysis of the average energy interactions defined by the third order tensor (4.16) defined previously. The structure of the interaction tensor **N** for a large POD-based model with $N = 75$, reconstructing more than 99% of the fluctuation kinetic energy, is reported in figure 5.12, showing the magnitude of the interactions for three slices for $i = 1, 10$ and 75, in panels (a), (b) and (c), respectively. It can be observed that all entries of the tensor **N** are generally nonzero, although the strength of the interactions varies across several orders of magnitude. This is a combined result of the projection coefficients tensor **Q** (shown later), whose entries are typically non zero, and of the complex spectral structure of the temporal coefficients $a_i(t)$, containing a blend of all the frequencies present in the system. The remarkable feature of figure 5.12 is that interactions are highly organised and there exists a subset of interactions that are more

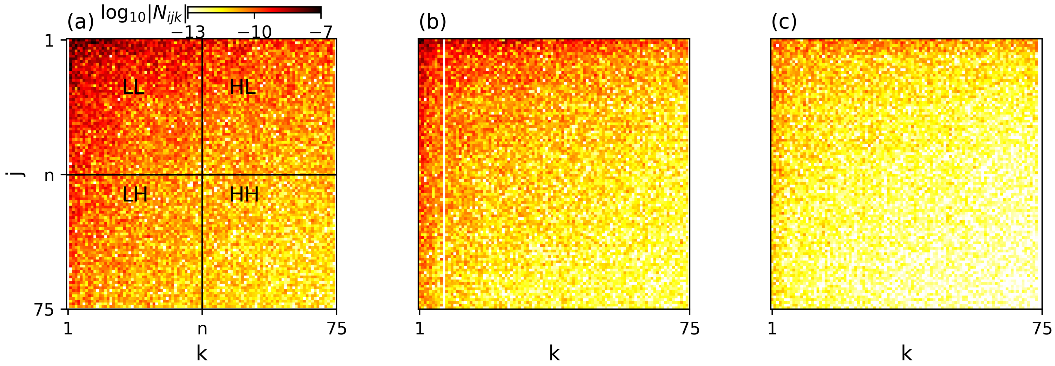

FIGURE 5.12: Magnitude of the average interaction tensor coefficients N_{ijk} for three POD modes across the spectrum, $i = 1, 10$ and 75 in panel (a), (b) and (c) respectively, for a model resolving 99% of the fluctuation kinetic energy.

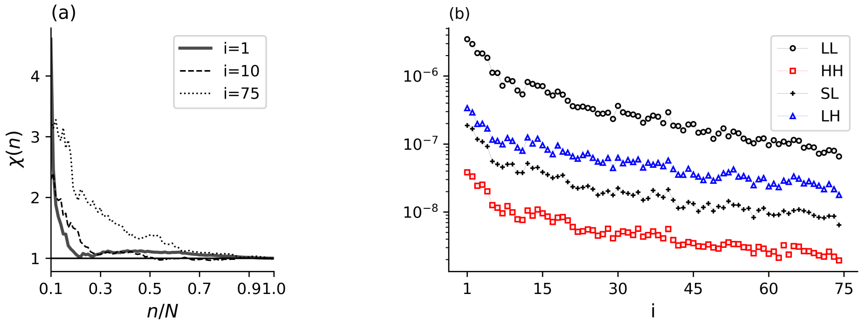

FIGURE 5.13: Panel (a): coefficient $\chi_i(n)$ as a function of the normalised cutoff n for the same three modes. Panel (b): sum of the entries of $\mathbf{N}$ the four regions defined in figure 5.12-(a) as a function of the modal index i.

active. Specifically, for any mode i, triadic interactions can be classified, as illustrated in panel (a), in four different categories by introducing a cutoff modal index n. The subset of interactions denoted as LL corresponds to nonlinear energy transfer involving pairs of low index modes, HL and LH denote interactions involving high-low/low-high index modes, while HH denotes the subset of interactions involving pairs of high index modes. Consequently, interactions LL, HH are local in modal space while interactions HL, HL are non-local. We observe that the areas corresponding to LL and HL/LH are the most active. If we map low/high modal indices to large/small scales, this result is in agreement with the picture of energy transfer between scales in homogeneous isotropic two-dimensional turbulence already observed by (Ohkitani, 1990; Laval et al., 1999), where the large scales interact with the small ones in a non-local fashion. In addition, interactions are not symmetric with respect to a swap of indices j, k. This feature can be quantified introducing the parameter

$$\chi_i(n) = \sum_{j=1}^{n} \sum_{k=1}^{N} N_{ijk} \Big/ \sum_{j=1}^{N} \sum_{k=1}^{n} N_{ijk}, \tag{5.12}$$

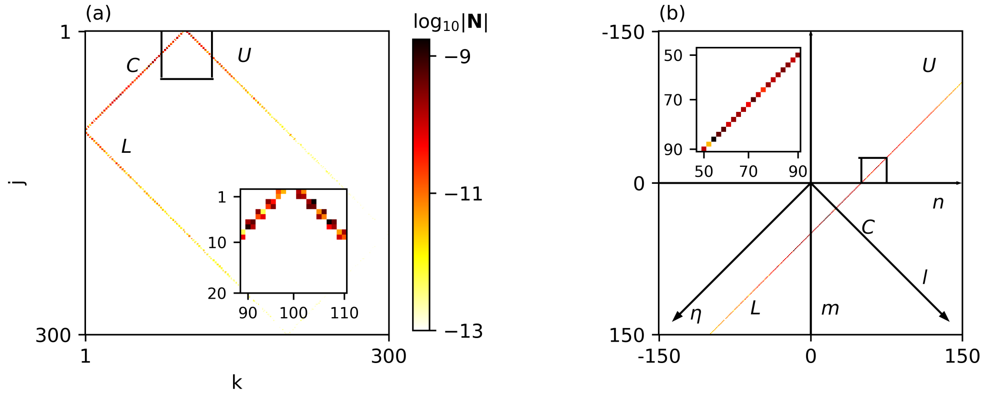

FIGURE 5.14: Panel (a): magnitude of the average interaction tensor **N** for $i = 100$, with the three characteristics branches, showing that active interactions come in 2×2 blocks corresponding to matching triads of modes. The small inset focuses on the interactions of branches C and U. Panel (b): magnitude of the average interaction tensor (5.15) where the three branches of panel (a) have been unfolded on a larger plane spanned by the coordinates l and η. The inset shows details of the interactions of the branch U in the plane $\eta - l$.

representing the relative dynamical importance of the subset of interactions LL + HL and LL + LH. Figure 5.13-(a) shows χ_i for $i = 1, 10$ and 75 as a function of the normalised cutoff n, as defined in figure 5.13-(a). The energy flow associated with the subset HL is up to four times stronger than the subset LH. This is a consequence of the asymmetry of the projection coefficients Q_{ijk}, which arises from the fact that the convective transport of structure $\boldsymbol{\phi}_k(\mathbf{x})$ operated by the structure $\boldsymbol{\phi}_j(\mathbf{x})$ is more intense when the modal structure $\boldsymbol{\phi}_j(\mathbf{x})$ describes large-scale flow features. To show this effect for all the different slices of N_{ijk} at i fixed we display in figure 5.13-(b) the sum of the absolute value of the entries of the four different regions of **N**, as defined in figure 5.12-(a). Here, we arbitrarily select the split at half of the spectrum ($n = 38$). Other choices are possible and do not change qualitatively the result. As expected, the intensity of the LL interactions is always much larger than the ones contained in the other regions. In addition, the energy transfer associated with the region LH is always larger than HL as consequence of the advective transport of the small structures by the large ones. More interestingly, the ratio LH/HL becomes larger as the modal index i increases. It is arguable, that this reflects the fact that as the flow structures become smaller the advective effect of the large structure becomes more important. This has also been observed more qualitatively in figure 5.13-(a).

We now consider energy analysis of a large, full-resolution DFT$_f$ model constructed from five partitions of thirty time units as discussed in section 5.2.2. The model is composed of all $N = 300$ modes, corresponding to 150 distinct frequencies. We perform modal decomposition and energy analysis on each partition separately, and then average the mean energy transfer rate tensor **N** over the five partitions. Figure 5.14-(a) shows the mean transfer rate distribution for mode $i = 100$. Energy interactions

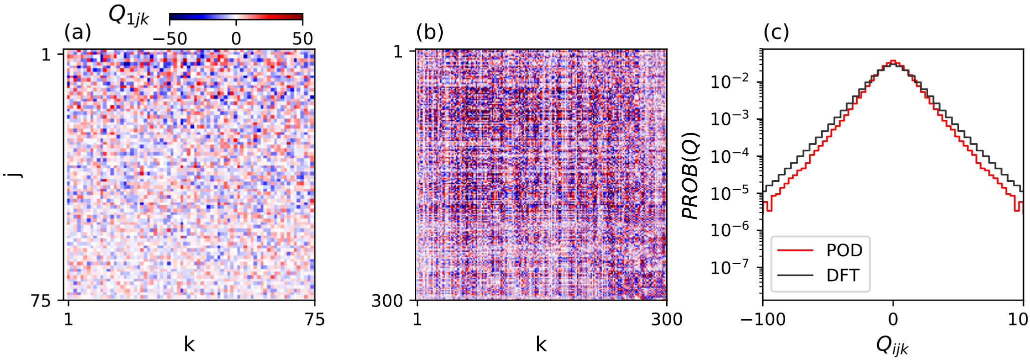

FIGURE 5.15: Maps of Q_{1jk} for Galerkin models constructed from the POD and the DFT_f decompositions, in panel (a) and (b), respectively. Panel (c) shows the probability distribution of all quadratic coefficients for these two models.

in the DFT model are very sparsely distributed on a thin horseshoe-shaped structure composed of three branches (denoted in the figure as L, C and U) of 2×2 blocks, and all other mean energy transfer rates interactions are identically zero. This pattern results from the joint effect of the oscillatory nature of the temporal coefficients and the quadratic nonlinearity of the system (4.6), which can only be satisfied by triads of modes having matching temporal wavenumbers. A less pronounced horseshoe-shaped distribution of the energy interactions has been previously observed in energy analysis of POD-based models of three-dimensional transitional boundary layer Rempfer and Fasel (1994a,b). These authors noticed that low-energy modal structures resemble Fourier modes in the spanwise direction (justified by the spanwise periodic domain) and thus coefficients Q_{ijk} and energy interactions are nonzero only for specific triads of modes. In the present case, this pattern is determined exclusively by the temporal coefficients as the tensor $\mathbf{Q}$ constructed from projection modes does not possess any sparsity structure and its coefficients have a similar statistical distribution to that obtained using the POD modes. This is illustrated in figure 5.15 showing maps of the first slice of the tensor $\mathbf{Q}$ of the largest Galerkin models considered here, constructed from the POD and the DFT decompositions, in panels (a) and (b), respectively. We observe that no underlying structure is present except for the asymmetry already observed in the energy analysis in figure 5.12. This property is confirmed in the probability distribution of the coefficients, shown in panel (c).

To facilitate the interpretation of the energy interaction pattern, we follow Rempfer and Fasel (1994b) and Arbabi and Mezić (2017) and define oscillatory modal structures

$$\mathbf{u}_l(t,\mathbf{x}) = a_{2l-1}(t)\boldsymbol{\phi}_{2l-1}(\mathbf{x}) + a_{2l}(t)\boldsymbol{\phi}_{2l}(\mathbf{x}), \tag{5.13}$$

numbered by the index l and tracing fluid motion at a single frequency on a two-dimensional subspace. Their modal energy is

$$e_l(t) = \frac{1}{2}\left(a_{2l-1}^2(t) + a_{2l}^2(t)\right) + a_{2l-1}(t)a_{2l}(t)(\boldsymbol{\phi}_{2l-1}(\mathbf{x}), \boldsymbol{\phi}_{2l}(\mathbf{x})). \tag{5.14}$$

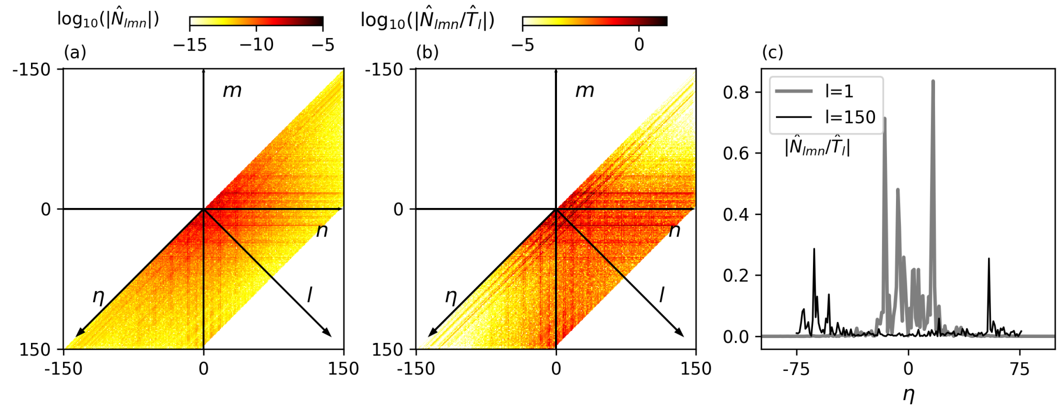

FIGURE 5.16: Panels (a) and (b): absolute and relative strength of the energy interactions between pairs of DFT modes for a model with $N = 300$ visualised on the plane m, n, with the additional coordinates l and η. Panels (c): relative energy interactions for the first and last mode pairs.

Numerical experiments show that, for large number of snapshots, pairs of modes $\boldsymbol{\phi}_{2l-1}(\mathbf{x})$ and $\boldsymbol{\phi}_{2l}(\mathbf{x})$ tend to be orthogonal. Hence, considering the evolution equation for the modal energy $e_l(t) \sim \frac{1}{2}(a^2_{2l-1}(t) + a^2_{2l}(t))$ leads to the condensed triadic interaction tensor $\hat{N}$ of size $(N/2, N/2, N/2)$ with entries

$$\hat{N}_{lmn} = \sum_{i=l}^{l+1} \sum_{j=m}^{m+1} \sum_{k=n}^{n+1} N_{ijk}, \tag{5.15}$$

lumping together the 2×2 blocks of interactions at matching triads of figure 5.14-(a). In addition, the three branches L, C and U can be unfolded and conveniently visualised on a two-dimensional plane spanned by the coordinate l, the modal structure index, and $\eta = m - n$, representing the distance in modal space between pairs of temporal wavenumbers. This unfolding process is shown in panel (b) of figure 5.14, and when repeated for all modal structures leads to the distribution shown in figure 5.16-(a). This diagram is exactly analogous to the ones shown in figures 5.2 and 5.3 with the only difference that in the case of the Kuramoto-Sivashinsky the complex nature of the Fourier modes introduces values of the wave-number k in $[-k_N, k_N]$. Conversely, the DFT modes are real with frequencies in $[\omega_1, \omega_N]$. In figure 5.16-(b), we report the average transfer rate $\hat{N}_{lmn}$ normalised with the total average transfer rate for each structure, the quantity

$$\hat{T}_l = \sum_{m=1}^{N/2} \sum_{n=1}^{N/2} \hat{N}_{lmn}, \quad l = 1, \dots, N/2, \tag{5.16}$$

to illustrate more clearly the relative strength of the interactions. In figure 5.16-(c), the normalised mean transfer rate for $l = 1$ and 150 is reported. Interactions between triads of pairs of DFT modes are organised in agreement with the physics of scale interactions previously discussed for POD models. In absolute terms, the most relevant interactions are clearly those located near the origin of the coordinates. These correspond to low-index modes where nonlinear interactions with other low-index modes dominate, while interactions with the high-index modes, for larger η, are less important. This suggests that a sparsification approach based on pruning the interactions involving the high-index modes, *i.e.* the small scales, would be effective. By contrast, for high-index modes, relevant energy interactions are organised in bands along the axes m and n and involve energy exchange between low-index modes and high-energy, high-index modes. This suggests that the dynamics of the small scales is driven primarily by non-local interactions with the largest structures of the flow and not by small-scale/small-scale quadratic interactions. The slight asymmetry visible in panel (c) arises from the structure of the coefficients tensor $\mathbf{Q}$ and has the same physical origin as that observed in figure 5.12 for the POD model.

5.2.4 Sparsification of POD-based models

We now apply the *a posteriori* sparsification to three POD-based models resolving 90%, 95% and 99% of the kinetic energy, respectively (see Table 5.1 for details). Because the size of the database matrix $\mathbf{\Theta}(\mathbf{A})$ grows quadratically with the number of modes, the number of possible interactions q can easily become larger than the number of available snapshots N_t, resulting in an underdetermined regression problem and overfitting. This is a well understood issue in data analysis and requires cross validation techniques to ensure the statistical reliability of the result (Friedman et al., 2008). In this work, we employed K$-$fold cross validation, using typically $\mathrm{K} = 10$. Briefly, the database is first divided into K folds. The model is trained using $\mathrm{K} - 1$ blocks and the reconstruction error ϵ of equation (4.23) is obtained from the fold that was left out. This procedure is iterated over all folds, obtaining the mean and the standard deviation of ϵ. Figures 5.17-(a,b,c) show the sparsification curves on the $\rho - \epsilon$ plane for the three POD models considered. The mean of ϵ across the folds is displayed as a thick black line, while the grey dashed line indicates plus or minus one standard deviation. These curves have been obtained by solving problem (4.22) using strategy S1 (*i.e.* $\gamma_i = \tilde{\gamma}$ is kept constant across the modes) and progressively increasing the regularisation weight. When low weights are used, dense systems (ρ near to 1) with good prediction accuracy are obtained, located in the right part of the panels. The opposite is true for large weights, identifying points in the left part of the graph characterised by low density and poor prediction accuracy. As postulated in the methodology section, the curves show a sweet spot at around $\rho \approx 0.2$, displaying a plateau for $\rho \gtrsim 0.2$, while the error ϵ grows quickly when $\rho \lesssim 0.2$. These results indicate that it is possible to prune about 80% of the quadratic interactions in model (4.6) without influencing the average prediction accuracy.

Overall, the mean reconstruction error decreases as the resolved energy increases, moving from panel (a) to panel (c), as more modes participate in capturing the dynamics of velocity fluctuations. In addition, larger models can be more effectively sparsified, as

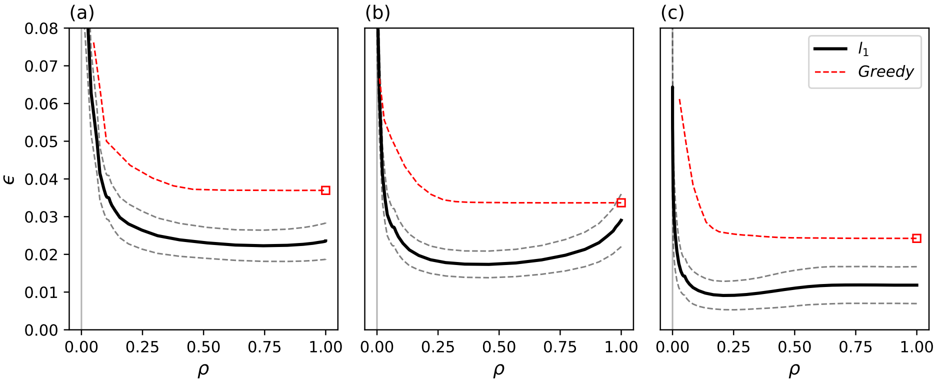

FIGURE 5.17: $\rho - \epsilon$ curves for three POD models resolving 90%, 95% and 99% of the kinetic energy in panels (a), (b) and (c), respectively. The black line represents the cross-validated error averaged over $K = 10$ folds. The dashed grey lines represent plus/minus one standard deviation of the cross-validated error calculated over the folds. The red dashed line shows the reconstruction error obtained with the greedy approach. The squares indicate the global reconstruction error of the Galerkin model obtained directly from projection.

the sparsification curve drops more rapidly. This results from the non-local structure of energy interactions shown in figure 5.12. When one additional low-energy mode is included, the number of relevant interactions to be retained in the model grows slower than $\mathcal{O}(N^2)$, i.e. all non-local interactions with the rest of the hierarchy denoted as LL, LH and HL in figure 5.12-(a) need to be considered ignoring the ones belonging to LH. Since the total number of possible interactions grows as $\mathcal{O}(N^3)$, larger models can be more effectively sparsified. This is conceptually in agreement with the observations of Taira et al. (2016) on the sparsification properties of discrete vortex models of increasing complexity. We also observe that the mean prediction error does not necessarily decrease monotonically when the density increases. This phenomenon is particularly visible for the model in panel (b) but all models reproduce the same behaviour. This is a symptom that the number of available snapshots (1500) is potentially not large enough for the number of coefficients ($q = N \times (N+1)/2 + N + 1 = 2926$ for the model in panel (c)) and overfitting would have occurred if no cross-validation had been performed.

5.2.5 Greeedy interaction selection

It can be observed that not all triadic interactions play the same role in the energy transfer of the system. In particular, concerning figure 5.12-(a), one could construct a model where the interactions labelled as HH are discarded. Hence, the simplest sparsification approach consists simply to construct a model where the retained interactions are manually selected by energetic considerations only without the need for any optimisation procedure. Here, we will refer to this sparsification technique as Greedy selection

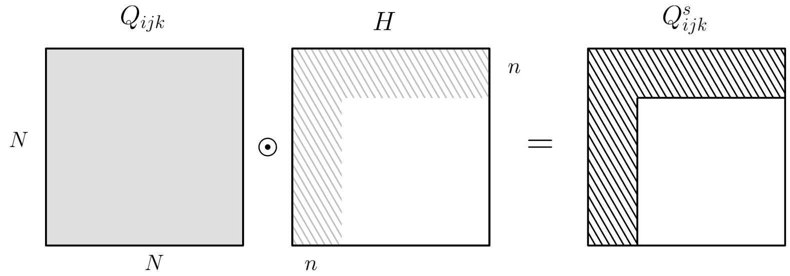

FIGURE 5.18: Greedy selection technique *via* Hadamard product between the tensor **Q** and the mask **H**. As a result, only the coefficients in the black shaded area are considered while the others are set to zero.

approach. More specifically, for this flow configuration, since in the tensor **N** the most relevant interactions are the ones belonging to the regions LL, LH and HL, the greedy sparsified model consists in retaining only the entries of **Q** corresponding to LL, LH and HL and setting to zero the complementary set. To practically perform the greedy sparsification we define a filter dependent on a parameter n, as shown in figure 5.12-(a), with a L-like shape similar to the contour of the contours of N_{ijk}. Thus, it is possible to generate different models with increasing values of densities varying the parameter n/N between 0 and 1. For this choice of the filter, the density can be defined as a function of n/N as

$$\rho = \frac{n}{N}\left(2 - \frac{n}{N}\right). \tag{5.17}$$

The filtering procedure is sketched in figure 5.18 and is performed simply by an entry by entry multiplication of the mask matrix defined as **H** and the i-th slice of the tensor **Q**. In tensorial notation this operation correspond to the Hadamard product defined as $\odot$. The sparse representation of the nonlinear interactions tensor **Q** is then defined as

$$Q^s_{ijk} = Q_{ijk} \odot H. \tag{5.18}$$

The red dashed line in figure 5.17 shows the $\rho - \epsilon$ curves is obtained by varying the ratio n/N. Then, the curves are constructed by computing the value of the acceleration $\dot{a}_i$ as right hand side of the expression (4.6) with the pruned third order tensor Q^s_{ijk}. The resulting $\dot{a}_i$ is then compared with the baseline DNS by the expression 4.23. In addition, the red square in the rightmost part of panels figure 5.17 displays the reconstruction error obtained with the dense/un-tuned Galerkin system. As expected from the iso-contours of the triadic interactions tensor N, the sparsification curves obtained with the greedy approach have a similar shape to those obtained with the l_1 regression. This is a direct consequence of the existence of a subset of most relevant energy interactions. However, the reconstruction error obtained from the greedy method is generally higher than that obtained by the l_1 regression, since the optimisation procedure involved in the l_1 approach modulates the strength of the remaining interactions by tuning the active quadratic coefficients, minimising the prediction error. As we show later, this difference will have a marked effect on the long-term temporal stability of the models.

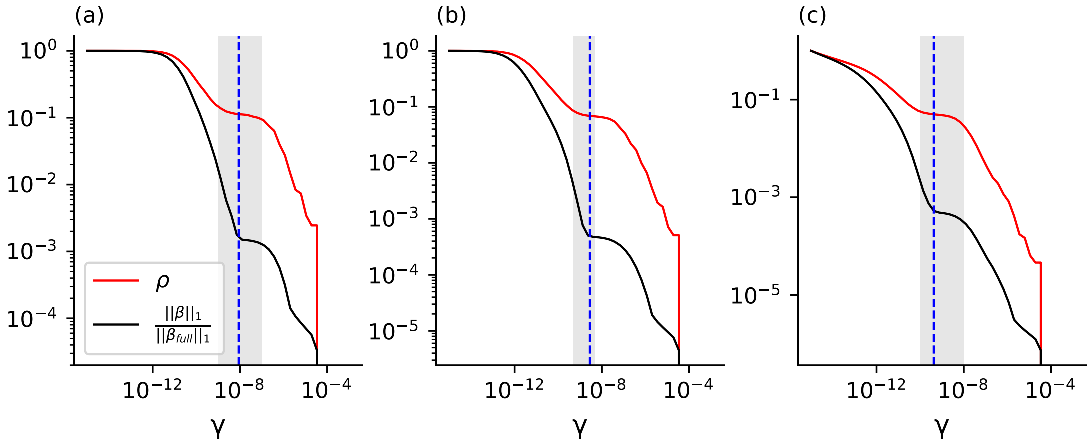

FIGURE 5.19: Value of the density ρ and of the pseudo-density ρ_{l1} as a function of the regularisation weight γ. Panel (a),(b) and (c) correspond to $e(n) = 0.9$,0.95 and 0.99, respectively. The blue vertical dashed lines indicates the optimal density according the criterion proposed here.

5.2.5.1 Selection of the optimal value of the regularisation weight

The sparsification procedure generates a whole family of models, dependent on the parameter γ_i, with different sparsity/reconstruction error trade offs that can be displayed on the $\rho - \epsilon$ plane as in figure 5.17. Models obtained with high or low values of γ_i must be ignored since the reconstruction error is too large (leftmost part of the diagram 5.17) or no sparsification at all is present (rightmost part of the diagram 5.17). However, for intermediate values of γ_i, there is a multitude of systems with similar values of the density/accuracy trade off among which the user can choose. Heuristically, a good graphical method to select the optimal value of γ_i is to pick a system in the proximity of the knee point of the curves 5.17. However, numerical experiments have shown that a manual choice of the model can be challenging since around the knee point the properties of the models can improve or deteriorate very quickly with just a small variation of γ_i. To address this issue, here, we propose a criterion to choose systematically the optimal regularisation weight $_i$ generating a sparse system with the best sparsity accuracy trade off. To our goal we first define a pseudo-density in terms of the l_1 norm as

$$\rho_{l1} = ||\boldsymbol{\beta}||_1 / ||\boldsymbol{\beta}_{full}||_1, \tag{5.19}$$

where $\boldsymbol{\beta}_{full}$ is the dense coefficient vector obtained by recasting the coefficients (4.7),(4.8) and (4.9) in vector form. Thus, we plot the trend of ρ and ρ_{l1} against the regularisation weight γ_i kept constant across the spectrum and indicated simply by γ. As expected, both curves are monotonically decreasing starting from one for very low γ (left) and approaching zero for large γ (right). Interestingly, we observe the presence of a plateau for both definitions of the densities for $\gamma \sim 10^{-8}$, in this area, the density ρ remains almost exactly constant while ρ_{l1} after hitting a sharp knee point keeps decreasing slowing.

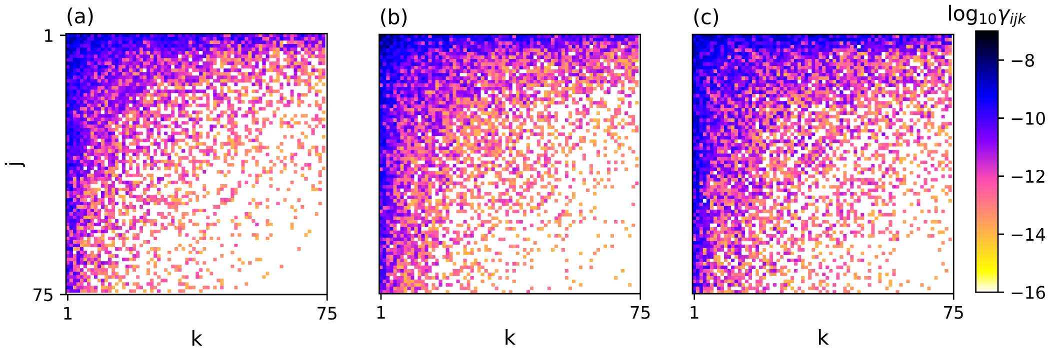

FIGURE 5.20: Distribution of the base ten logarithm of γ_{ijk} for $i = 1, 10$ and 75, in panels (a), (b) and (c), respectively, for the POD model resolving 99% of the fluctuation kinetic energy.

Our interpretation is that, for these intervals of the regularisation weight (grey shaded area in figure 5.19) no coefficients are eliminated but their values are being tuned. We argue that the optimal regularisation weight corresponds to the knee point in the ρ_{l1} curve (indicated as a vertical dashed blue line) where all the un-necessary entries of the matrix coefficients are easily eliminated keeping only the coefficients corresponding to relevant energy interactions. A further increase of γ results in the elimination or too large modification of the retained coefficients, including relevant interactions as well, leading to a system with poor prediction accuracy. Similarly to what is shown in figure 5.17 we observe that the values of γ corresponding to the optimal density decreases monotonically as the energy resolution of the model increases. This behaviour has been already observed in a more qualitative fashion in figure 5.17 this is a consequence of the increasing complexity of the system due to the structure of the triadic interaction tensor **N** where the number of irrelevant interactions grows faster in relative terms with respect to the important ones making the system more easily sparsifiable.

5.2.6 Energy interactions identified by the regression and conservation properties

To visualise the sparsity pattern identified by the regression as the regularisation weight in equation (4.22) is increased (constant for all modes in strategy S1), we introduce the tensor γ with entries γ_{ijk} defined as the value of the regularisation weight at which the corresponding coefficient Q_{ijk} is shrunk to zero by the LASSO. Figures 5.20-(a,b,c) show three slices of γ for modes $i = 1, 10$ and 75, respectively, for the largest POD model considered, capturing 99% of the total fluctuation kinetic energy. The first interactions to disappear are the small-scale/small-scale interactions, region HH. Increasing the penalisation, interactions that are local in modal space are progressively pruned, leaving

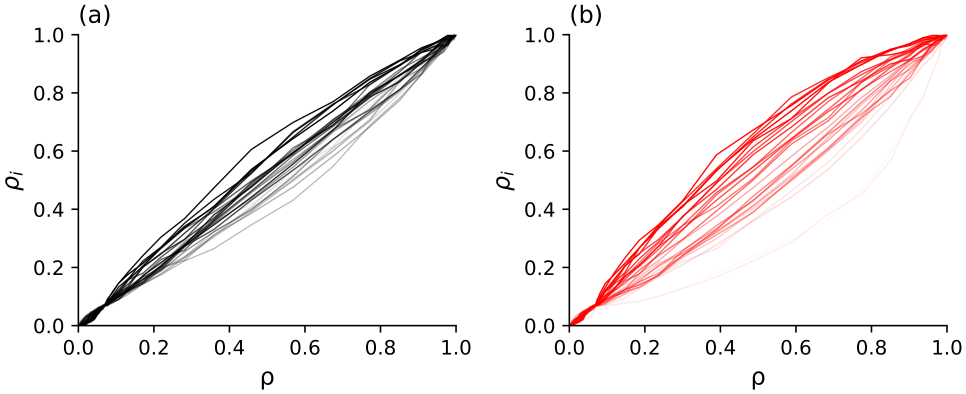

FIGURE 5.21: Value of the modal density ρ_i against the global density ρ. The intensity of the colour is proportional to the modal index i (low intensities correspond to low values of i and the opposite is true for high intensities). The sparsification approaches S1 and S2 are displayed in panel (a) and (b), respectively.

only non-local interactions involving triadic exchanges with the low-index modes for large penalisations. Interestingly, this pattern does not change qualitatively nor quantitatively as the modal index i increases.

To better visualise how the sparsification is distributed across the modes the value of the modal density ρ_i is plotted against the global density ρ in figure 5.21. The approaches S1 and S2 are shown in panels (a) and (b), respectively. The opacity displays the value of the modal index. Opaque lines correspond to high index low energy modes while transparent lines correspond to low index high energy modes. It can be observed that, for a given value of the global density ρ, a comparable number of interactions is retained across the hierarchy of modes and each modal equation is sparsified by a similar amount, this is true for both approaches S1 and S2. Hence, the sparsification has not produced mode truncation, which would have occurred if all coefficients of some low-energy modes had been shrunk to zero by the LASSO. Truncation would result in ρ_i dropping to zero while being the global ρ different from zero. In addition, it can be observed that the low index modes (more transparent) are more effectively sparsified both by the S1 and S2 approach as consequence of the structure of the triadic interaction tensor shown in figures 5.12 and 5.13. Since, as shown in 5.13-(a), low indexes modes span larger energy ranges with respect to the modes in the rear part of the spectrum, as it can be seen in figure 5.13-(c).

In figure 5.22-(a,b,c) the base ten logarithm of the mean energy interaction tensor $\mathbf{N}^s_{ijk}$ computed as in (4.16) with the sparse coefficient tensor $\mathbf{Q}^s$ is shown for the same three modal indices as in figure 5.20. Data refers to sparse models with $\rho = 0.3$ identified with the methodology explained in the previous section and located nearby the sweet spot of the curves in figure 5.17-(c). It can be observed that the sparsified model has a pattern of interactions resembling that of the dense model in figure 5.12. However, weak interactions and the associated flow physics have been pruned. It is also clear

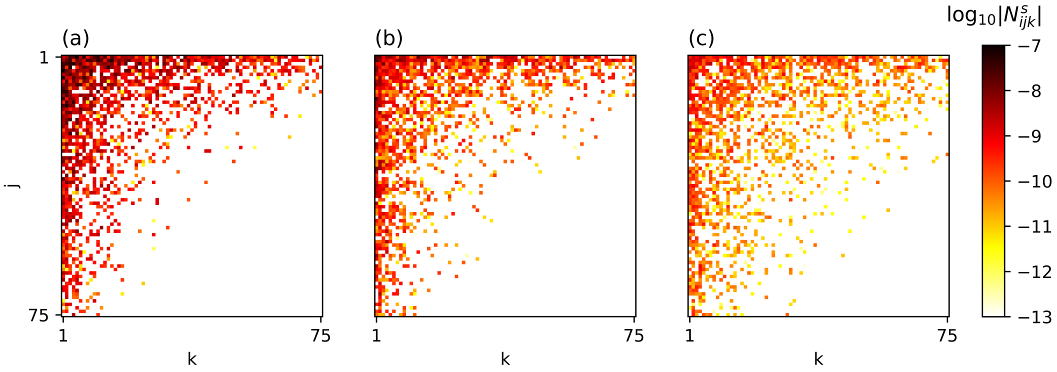

FIGURE 5.22: Base ten logarithm of the sparsified interaction tensor N^s_{ijk} for $i = 1, 10$ and 75 in panels (a), (b) and (c), respectively, for a POD model model resolving 99% of the fluctuation kinetic energy and $\rho = 0.3$.

that the asymmetry of the interaction pattern observed in figure 5.12 and the physical mechanism that originates is lost in the regression since the database matrix $\mathbf{\Theta}(\mathbf{A})$ does not include any permutation of the modal indexes. Consequently, the interaction pattern in figure 5.22 is now symmetric with respect to a swap of the indices j, k.

Despite the aggressive pruning, the sparse models reproduce fairly accurately the overall structure of the intermodal energy budgets. In the present flow configuration, the convective nonlinearity is energy conserving and Galerkin models should obey the relation $\sum_{i=1}^{N} T_i = 0$ as $N \to \infty$, with $T_i = \sum_{j=1}^{N} \sum_{k=1}^{N} N_{ijk}$ the time averaged energy transfer rate to/from mode i. For finite-dimensional approximations, this property is not satisfied exactly and the residual of the summation can be taken as a measure of the overall energy balance. Figure 5.23-(a) shows such residual for the l_1 sparsified models (empty circles) and for the models obtained with the greedy approach (empty squares), as a function of the density. The residual is normalised by the root mean square value of the rate of change of the integral fluctuation energy. Note that the greedy model at $\rho = 1$ is the model obtained directly from projection. It can be observed that the energy conservation error is relatively small, in the order of $10^{-3} \div 10^{-4}$. For large densities, it is larger than that of the projection model, because the regression tunes model coefficients to minimise the mean square error on the modal accelerations and does not enforce this physical constraint directly. The energy conservation residual decreases for sparser models and is ten times smaller than the projection model, owing to the lower number of active coefficients that participate in the regression. Figure 5.23-(b) shows the distribution of the time averaged energy transfer rate associated to mode i for the l_1 sparsified model at $\rho = 0.3$ (red crosses) and the dense model obtained from projection (empty circles). Data is reported every second mode. For the projection model, the net

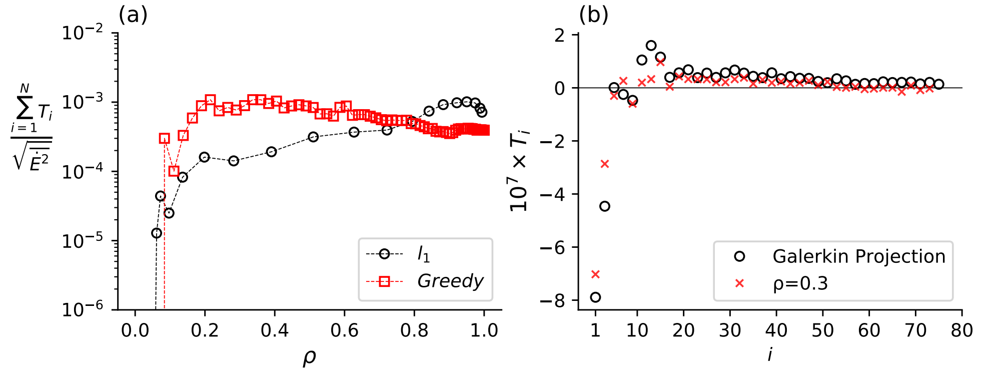

FIGURE 5.23: Panel (a): the normalised energy conservation error as a function of the density ρ, for the greedy sparsification approach (empty squares) and the l_1 based approach (empty circles). Panel (b): the total net energy transfer rate T_i as a function of the modal index i for two POD models resolving 99% of the kinetic energy, with coefficients identified from projection and for a l_1 sparse model. One every two data points is reported.

energy transfer is negative for the first few modes and changes sign at $i \sim 10$. Physically, this trend suggests that the first few modes extract energy from the mean flow and feed the dissipative high-index modes via triadic interactions. The l_1 sparsified model correctly reproduces this global trend, even though no constraints have been introduced (Loiseau and Brunton, 2018). Generally speaking, as discussed by (Loiseau and Brunton, 2018) and (Taira et al., 2016), machine learning techniques, purely based on data, do not automatically include physical invariants. The evidence shown here and later on in the section dedicated to the *a priori* approach indicates that physically consistent results can be obtained if constraints are imposed (i.e. time stability constraint for example). However, it can be argued that data-driven techniques relying on optimisation ideas, such as the LASSO, can naturally reproduce invariants and conservation properties embedded in the data to a level defined by noise levels. Since the optimisation procedure tends to naturally reproduce the invariants contained in the original data. For instance, Taira et al. (2016) used network-theoretic ideas to sparsify connections in a discrete vortex model and observed that sparsification conserves the invariants of discrete vortex dynamics. The critical aspect is therefore the choice of sensible physics preserving constraints and, at the same time, allowing enough degrees of freedom to successfully perform the fitting procedure.

Lastly, it is important to underline that, from figure 5.17 seems that the models with the best trade off between sparsity and accuracy are located near values of $\rho = 0.2$ corresponding to models located nearby the knee point. This analysis shows that models characterised by a slighter higher value of ρ, identified by the approach defined in figure 5.19, reproduce better the energy transfer inside the model. This result underlines, once more, the difficulties in the choice of γ without *a posteriori* evaluation of the performances of the sparse system obtained by temporal integration.

5.2.7 Regression set-up for pure frequency modes

Proper Orthogonal Decomposition is a popular and well tested modal decomposition to generate reduced order models. However, this is not the only choice when it comes to the modal analysis of flows. An interesting class of modal decompositions, particularly suitable for statistical stationary flows, are the ones generating spatial structures whose temporal coefficients oscillate at a fixed frequency. Several formulation have been proposed such as the Spectral POD (see Towne et al. (2018); Sieber et al. (2016)) and the Dynamic Mode Decomposition DMD (Schmid (2010)) particularly suitable for fully developed flows. In addition, one of the targets of this work is to show how to sparsify reduced order models preserving the physical features of the original flow regardless of the set of basis functions used to represent the flow. To prove that the sparse system is physics preserving we generate a Galerkin-based model using pure frequency modes obtained by Discrete Fourier Transform (DFT). In this section, we show the precautions needed to be taken into account before solving problem (4.22) when pure frequency modes are considered.

5.2.8 Ill-posedness of the problem

As an illustrative example, we consider a small-sized model of the lid driven cavity flow at Reynolds 2×10^4 constructed with $N = 26$ modes and perform sparsification as discussed in the methodology section, with a relatively small regularisation weight ($\gamma = 10^{-14}$) kept constant across the frequency spectrum (strategy S1). The first key result is that all the entries of the constant and quadratic coefficient tensors **C** and **Q** are immediately set to zero, while the linear tensor **L** has a characteristic bidiagonal structure, shown in figure 5.24-(a). The system identified by the regression is equivalent to a set of $N/2$ decoupled linear oscillators in the form

$$\begin{bmatrix} \dot{a}_{2l-1} \\ \dot{a}_{2l} \end{bmatrix} = \omega_l \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \begin{bmatrix} a_{2l-1} \\ a_{2l} \end{bmatrix} \qquad l = 1, ..., N/2, \tag{5.20}$$

coupling pairs of temporal coefficient oscillating at the same angular frequency $\omega_l = 2\pi l/T$, with T being the observation time. The eigendecomposition of the tensor **L** is trivial. Eigenvalues are all imaginary and come in pairs that are integer multiples of the fundamental frequency $\omega_1 = 2\pi/T$. While this result is consistent with recent ideas on Koopman operator theory (Mezić, 2013), where nonlinear dynamics are modelled with a linear system of larger dimension, all information on nonlinear energetic interactions has been lost in the process since the nonlinear part of the system has been completely eliminated by the regression. This result is due to the fact that, when temporal coefficients are sine/cosine pairs, there is a column of $\boldsymbol{\Theta}(\mathbf{A})$ that is exactly parallel to the target $\dot{\mathbf{A}}_i$, since time differentiation is equivalent to a permutation of sine/cosine pairs.

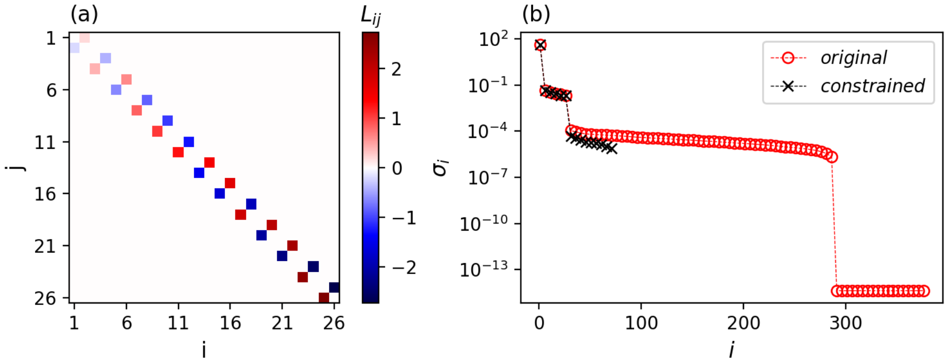

FIGURE 5.24: Panel (a): entries of the linear tensor **L** identified by the unconstrained regression. Panel (b): singular values σ_i of the full database matrix $\mathbf{\Theta}(\mathbf{A})$ of equation (4.20) (red circles), and of the reduced matrix (black crosses) obtained by keeping only the subset of columns corresponding to active interactions on the three active branches of **N**. One every five singular values is shown for clarity.

As pointed out in Brunton et al. (2019) incorporating and enforcing known flow physics is a challenge and opportunity for machine learning algorithms. To address this first aspect, we introduce a physically-motivated approach based on considerations of the time averaged energy budget of system (4.12) where only the diagonal elements of the linear term participates in the mean power budget. Hence, for the sparsification of DFT-based models we use a modified database matrix that only contains the column associated to the diagonal part of the linear term. The second aspect is that for DFT models the database matrix $\mathbf{\Theta}(\mathbf{A})$ is not full rank. This is direct consequence to the particular oscillatory structure of the temporal DFT modes $a_i(t)$. More specifically, since $a_i(t)$ come into pairs of sin, cos functions and the database matrix $\mathbf{\Theta}(\mathbf{A})$ contains only linear and quadratic combinations of $a_i(t)$. This results in some of the columns of this matrix are linearly dependent. In this case, the LASSO is known to select one column at random (according to the particular ordering of the columns) and sets to zero regression coefficients of the other linearly dependent columns (Tibshirani, 2013; Hastie et al., 2015). Machine learning techniques often come without guarantees for robustness (Brunton et al., 2019), implying that physical insight obtained with these tools might be questionable. To avoid this problem, we constructed a reduced database matrix $\mathbf{\Theta}(\mathbf{A})$ containing only columns corresponding to the entries of **Q** corresponding to the interactions of **N** we know being non zero due to the relation between triad of frequencies such that $i + j = k$. This features has been observed approximately by Rempfer and Fasel (1994a) and it is satisfied exactly in this case. Thus, relevant interactions in the constrained problem can be identified via regularised regression (4.22). The reduced database matrix is full rank, as can be seen in panel (b) of figure 5.24, showing the singular values of the full database matrix defined by equation (4.20) and of the reduced matrix. The important consequence is that the solution of the LASSO problem (4.22) is unique (Tibshirani, 1996), and can be thus compared with the available physical

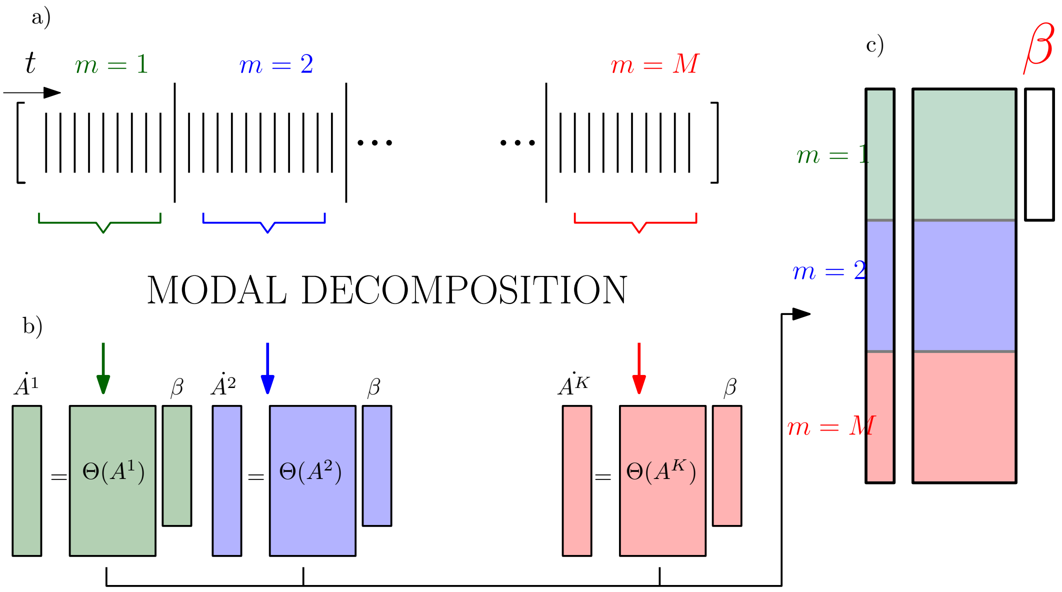

FIGURE 5.25: Panel (a): subdivision of the time signal into M partitions. Panel (b): generation of one LASSO problem for each partition. Panel (c): staking of all M problems to generate one LASSO problem (4.22) for a longer time interval.

knowledge of scale interactions in turbulent flows. This reduction is not strictly necessary, since the l_1 regression identifies this pattern anyway for fairly small regularisation weights. However, this has the advantage that the computational complexity of sparsifying the entire Galerkin model only grows as $\mathcal{O}(N^2)$ instead of $\mathcal{O}(N^3)$, as for POD models, because the reduced database matrix contains a number of interactions equal to $q = 2(N+1)$ at most. More importantly, because of the greatly reduced number of free coefficients cross-validation techniques to avoid over-fitting become unnecessary.

5.2.9 Ambiguity in the definition of the problem

The third aspect to be taken into consideration when considering DFT-based models is that the number of modes and, consequently, the properties of the spatial and temporal modes $a_i(t)$, is not uniquely defined by the energy resolution (as for the POD) but depends on the observation time T as well. Long observation times would be beneficial to reach statistical significance in the solution of (4.22) but would result in low-energy/low-frequency modes that do not contribute significantly to the overall dynamics adding complexity to the problem. To tackle this issue a methodology based on the partitioning of the dataset shown in figure 5.25 has been developed. First, we divided the original dataset into M partitions, as shown in figure 5.25-(a), and performed DFT for each of them separately obtaining M different sets of DTF modes. For each partition we compute the modal acceleration $\dot{\mathbf{A}}_i^m$ and we assemble the corresponding database matrix $\mathbf{\Theta}(\mathbf{A}^m)$. Thus, we assemble M different problems (4.22) one for each partition, as shown in figure 5.25-(b). Lastly, we stacked vertically the modal acceleration matrices and the reduced database matrices from the partitions and solved (4.22)

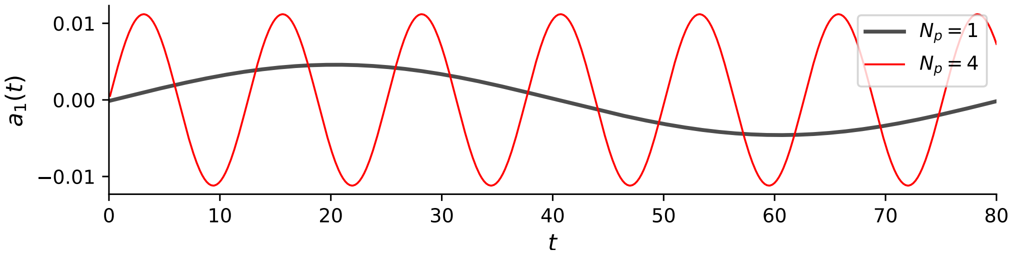

FIGURE 5.26: First temporal DFT mode with the observation time is equal to the available time span $T = 80$ in black. First temporal DFT mode when the observation time is equal to $T = 40$ and the dataset is divided into four partitions, in red.

for a common coefficients vector $\boldsymbol{\beta}_i$ as displayed in figure 5.25-(c). The result of this methodology is shown in figure 5.26 showing in black the first temporal mode (with the lower frequency) $a_1(t)$ of the DFT decomposition obtained with an observation time equal to the available timespan $T = 80$ time units. Conversely, in red the first DFT mode obtained with a subdivision of the dataset in four partitions ($T = 20$) is shown. It can be observed how the partitioning helps to generate low indexes modes already containing a considerable amount of kinetic energy and to get rid of low frequency and low energy physically meaningless modes.

5.2.10 Sparsification of DFT-based models

We now move to the sparsification of DFT-based models. We focus primarily on the structure of energy interactions identified by the regression and leave long-term temporal stability considerations aside. In fact, the DFT produces modal structures by assuming *a priori* their temporal behaviour, i.e. harmonic motion, and the meaning of a time-domain analysis is thus conceptually unclear.

We introduce the modified density ρ_{DFT} spanning the range $[0, 1]$ and representing the number of active coefficients with respect to the total number of active interactions on the three branches of figure 5.14. For large models, the approximation $\rho_{\text{DFT}} \approx 2/3\rho N$ can be used. In figure 5.27-(a), sparsification curves for three models obtained with observation times $T = 10$, 30 and 50 (with $M = 15$, 5 and 3 partitions of the full dataset, respectively), at full energy resolution, are reported. Strategy S1, where the regularisation weight is maintained constant for all modes is used.

Similarly to what observed for the POD-based Galerkin models, the reconstruction error decreases monotonically with the observation time. This is a consequence of the larger number of frequencies that interact quadratically to reconstruct the original DNS acceleration data. For the larger model obtained at $T = 50$, 70% of the triadic interactions can be pruned with no major effects on the overall prediction error. If the full coefficient tensor $\mathbf{Q}$ is considered, this correspond to a remarkably low density of 0.0015.

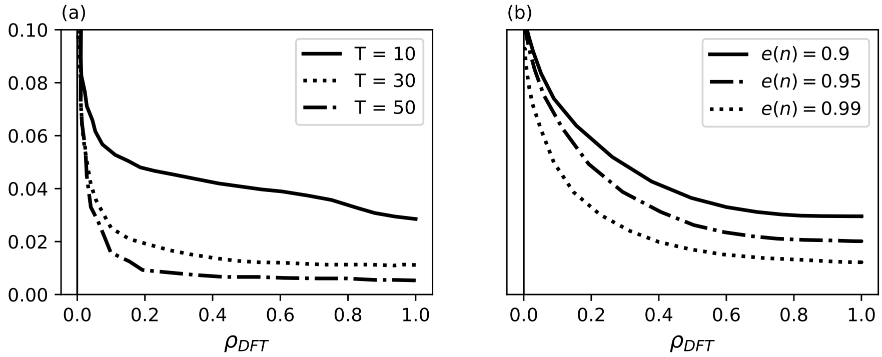

FIGURE 5.27: Panel (a): sparsification curves for models obtained by three different observation times and resolving 100% of the kinetic energy. Panel (b): sparsification performed with $T = 30$ with three different energy resolutions $e(n)$.

Figure 5.27-(b) shows the sparsification curves for models obtained with observation time $T = 30$, for three different energy resolutions, $e(n) = 0.9$, 0.95 and 0.99. Interestingly, we notice that the curves do not present a large plateau for high densities as opposed to the full resolution mode shown in panel (a) and the POD sparsification curves of figure 5.17. This is the combined effect of the dramatic decrease in the number of modes at lower energy resolutions (see table 5.1) and the inherent efficient description of energy interactions in DFT-based models compared to POD.

We now compare strategies S1 and S2 on the full resolution model obtained with observation time $T = 30$. Results of this analysis are reported in figure 5.28. The top/bottom panels are obtained with the strategy S1/S2, respectively. Panel (a) shows the tensor $\hat{\gamma}$, obtained by processing and visualising the full tensor γ using the same technique utilised for the interaction tensor $\hat{\mathbf{N}}$ in figure 5.14. Panel (b) shows the density of individual ordinary differential equations for a selected number of modal structures as a function of the overall model density ρ_{DFT}, while the sparsified interaction tensor $\hat{\mathbf{N}}^s$ for $\rho_{\text{DFT}} = 0.7$ is shown in panel (c). When the regularisation weight is maintained constant, the sparsification pattern emerging from the tensor $\hat{\gamma}$ follows the distribution of the mean energy transfer rate of figure 5.16-(a). In particular, despite the signature of non-locality is still visible in the pattern, the sparsification is highly skewed across the spectrum because the equations for high-frequency modes are excessively sparsified for moderate penalisations as opposed to those of low-frequency, high-energy modes. This behaviour is better seen in the individual density curves in panel (b). Specifically, the density ρ_l of the last mode pair ($l = 150$) drops quite pronouncedly to much lower density than average at $\rho_{\text{DFT}} \approx 0.5$. Panel (d) shows the sparsification pattern obtained with the second strategy. We observe that, in this case, the interactions are retained according to their relative strength producing a sparsification pattern that follows the relative energy transfer rate reported in figure 5.16-(b). This results in a more balanced sparsification across the spectrum, where the modal density ρ_l decreases more uniformly for all modes as the global density is decreased, as shown in panel (e). The mean

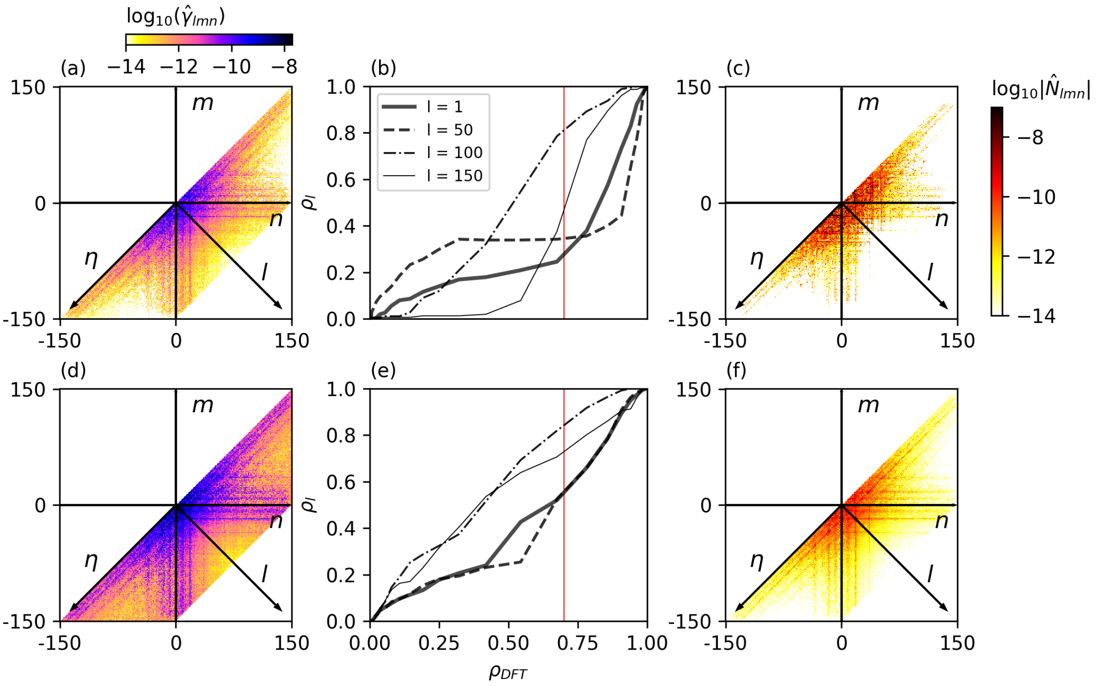

FIGURE 5.28: Top panels: strategy S1; bottom panels: strategy S2. Panels (a) and (d) show the distribution of $\hat{\gamma}$. Panels (b) and (e) show the trend of the modal density ρ_l against the global density ρ_{DFT} for four different modes in different parts of the spectrum. Panels (c) and (f) show the energy interaction tensor N^s of the sparsified system.

energy transfer rate of the models sparsified using the two strategies, with $\rho_{\mathrm{DFT}} = 0.7$, is reported in panels (c) and (f). Globally, the structure and intensity of energy interactions is preserved by the LASSO, although strategy S1 has more aggressively sparsified the high-index modes and truncated the equations of the last five pairs of modes. Although not shown here, all DFT models, regardless the strategy, have similar energy conservation properties as for POD modes, as illustrated in figure 5.23.

As a final remark, we have observed in sparsification of larger DFT models that, although the LASSO is able to successfully identify the dominant subset of energy interactions, the complexity of the optimisation problem makes an accurate reconstruction of the numerical values of the system coefficients challenging. This is due to the spectral properties of the database matrix $\mathbf{\Theta}(\mathbf{A})$ which deteriorate as the number of modes considered grows (Cordier et al., 2010). A potential solution to this issue would be to use elastic-net regression (Friedman et al., 2008) which combines an l_1 term with an l_2 (Tikhonov) penalisation. This would provide a better trade-off between sparsification and stability of the reconstructed coefficients.

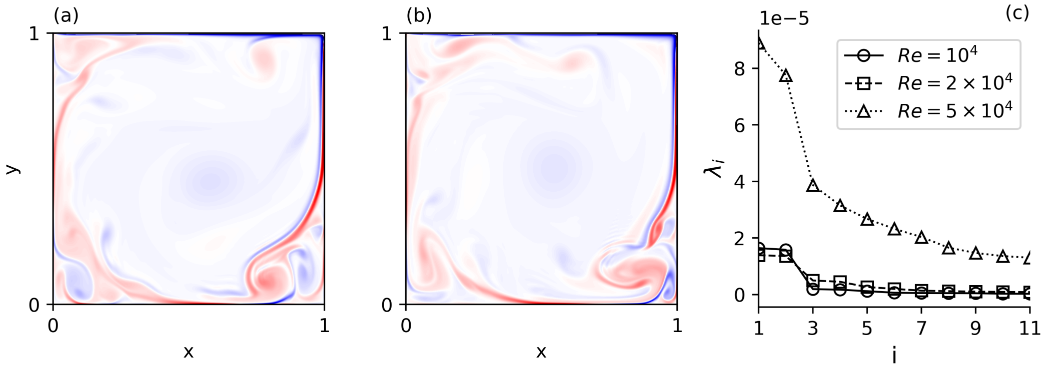

FIGURE 5.29: Two vorticity snapshots of the cavity flow at $Re = 5 \times 10^4$ in panels (a) and (b), respectively. Panel-(c): first 10 POD eigenvalues λ_i for the three Reynolds numbers considered.

5.3 Effects of the Reynolds number

To study the effects of the Reynolds number on the l_1 sparsification framework other two flow configurations, at $Re = 10^4$ and $Re = 5 \times 10^4$, of the lid driven cavity flow are considered. Thus, these additional test cases are integrated with the results obtained with the previously analysed flow at $Re = 2 \times 10^4$. While the flow dynamics for $Re = 10^4$ is similar to the one of $Re = 2 \times 10^4$, the flow-field obtained with the highest Reynolds number displays different features. These are shown in figure 5.29 panels (a) and (b) where two vorticity snapshots for the cavity flow at $Re = 5 \times 10^4$ are displayed. It can be observed the presence of a strong recirculation vortex in the lower right corner that sheds erratically breaking the shear layer while being advected downstream by the mean flow. This phenomenon is accompanied by strong, quasi-periodic bursts in the turbulent kinetic energy. This feature makes the statistical description of the phenomenon challenging, requiring a large number of snapshots to reach convergence of the statistics. Figure 5.29-(c) displays the first 10 POD eigenvalues for the three Reynolds numbers considered. It can be observed that for the two lowest values of the Reynolds numbers the first two eigenvalues come as a pair with the same energy. This behaviour has been previously previously observed for the cavity flow at $Re = 2 \times 10^4$ and it is due to the presence of the strong quasi-periodic shear layer as shown in figure 5.9 and 5.11 for the two lowest Reynolds numbers analysed. Conversely, for $Re = 5 \times 10^4$ this feature is not present since the vortex detachment makes the shear layer unstable, eventually breaking it and no wavelike description of the phenomenon is possible. Table 5.2 shows the number of POD modes needed to reach the desired values of energy resolution for the three different Reynolds numbers considered. It can be observed that the number of modes does not increase linearly neither with the Reynolds number nor with the energy resolution leading to an impractically large number of modes to recover large amounts of kinetic energy for the highest Reynolds number considered. This result shows the technical difficulties that

$e(n)$	0.9	0.95	0.99
$Re = 10^4$	8	12	36
$Re = 2 \times 10^4$	20	35	75
$Re = 5 \times 10^4$	60	112	290

TABLE 5.2: Number of modes required as a function of Reynolds number and energy resolution.

need to be tackled while analysing highly resolved models of high Reynolds number flows.

5.3.1 Sparsification curves

Figure 5.30 summarises the results of the sparsification procedure applied to the models for the two additional Reynolds numbers alongside the model at $Re = 2 \times 10^4$ already considered in figure 5.17. In each panel, the value of the mean cross validated error ϵ is displayed against the system density ρ. For the sake of clarity each Reynolds numbers is plotted with different line-styles: continuous,dashed and dotted for $Re = 10^4, 2 \times 10^4$ and 5×10^4, respectively. The value of the resolved kinetic energy increases from left to right being $e(n) = 0.9, 0.95$ and 0.99 in panel (a,b,c), respectively. Similarly to the previous analysis, when low regularisation weights are used (points in the right part of the graphs), dense systems with low prediction accuracy are obtained. The opposite is true for large regularisation weights, identifying points in the left part of the graphs. As postulated, since a set of coefficients in the tensor $\mathbf{N}$ is predominant, the curves present an initial plateau for high densities where it is possible to remove coefficients from the system (4.6) without significantly affecting the reconstruction error ϵ. Generally, the sparsification curves in figure 5.30 show two common emerging trends. First, the value of the reconstruction error ϵ in correspondence of the plateau decreases monotonically as the energy resolution increases, since more modes participates in capturing the dynamics of the fluctuations. Second, and more importantly, the optimal density ρ_{opt} decreases as the Reynolds number increases. The value of the optimal density ρ_{opt}, for each model taken into consideration, computed with the procedure explained in figure 5.19 is displayed in figure 5.31-(a). As expected, it can be observed that ρ_{opt} decreases monotonically both as the energy resolution increases (moving vertically in the graph) and as the Reynolds number increases (moving horizontally). This confirms what already partially observed for $Re = 2 \times 10^4$, *i.e.* that more complex models can be more efficiently sparsified. This can also be observed qualitatively in figure 5.30. Crucially, this features has direct effect on the computational cost of the time integration of the models themselves. It is important to underline that the actual computational cost, *i.e.* wall-time taken to advance the model by one time unit, highly depends on a number of implementation-dependent factors and code optimisations. For instance, for dense models, one could store the quadratic tensor $\mathbf{Q}$ in slices for each i and use BLAS

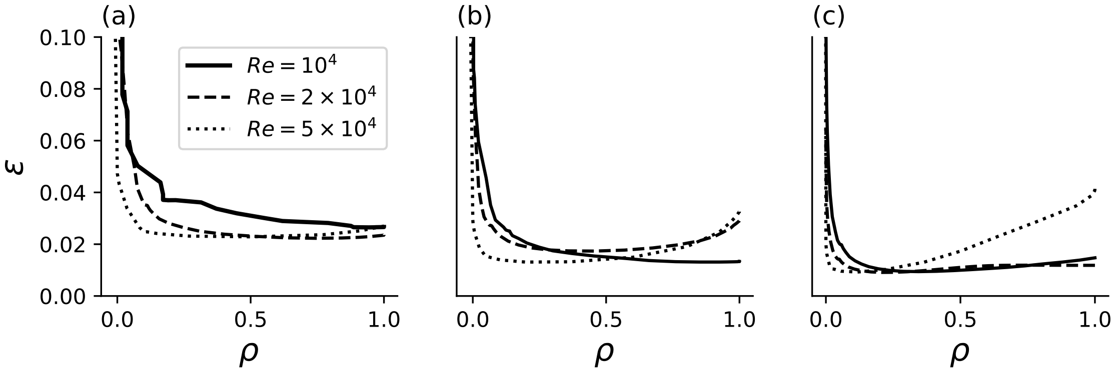

FIGURE 5.30: Mean of the cross–validated reconstruction error against the density of the system, ρ: in (a), $e(n) = 0.90$; in (b), $e(n) = 0.95$; and in (c), $e(n) = 0.99$.

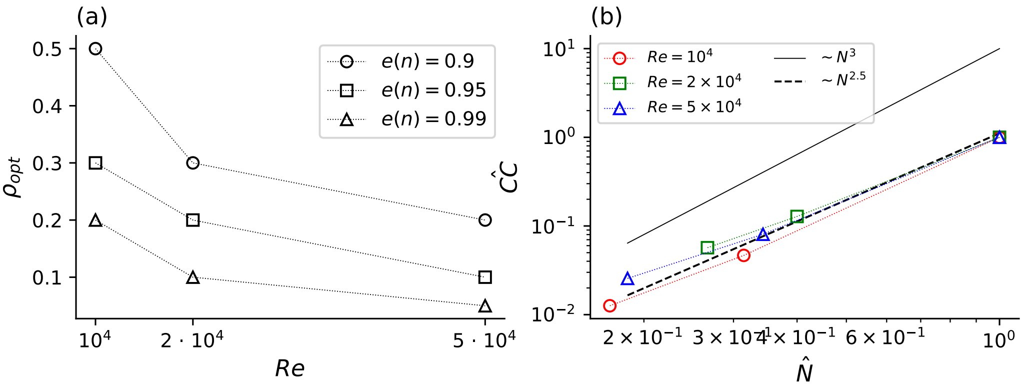

FIGURE 5.31: Panel-(a): optimal density against the Reynolds number for the systems considered in figure 5.30. Panel-(b): normalised computational cost against normalised number of modes for the 9 systems considered as red dots. The panel displays the asymptotic trend as continuous line and the linear interpolation of the red dots as dashed lines.

matrix-vector routines. For sparse models, one could use sparse-matrix techniques to take advantage of the sparsity. In this analysis, to rule out every aspect related to the implementation we define the computational cost as the number of operations needed to evaluate the quadratic term in (4.22). Thus, for a dense system composed by N modes, computational costs must scale with N^3. On the other hand, for sparse systems, computational costs scales with the number of active coefficients, *i.e.* ρN^3. Figure 5.31-(b) shows the computational cost of each model considered against the number of modes needed to reconstruct the target amount of turbulent kinetic energy. The x and y axes are scaled to have values in the interval $(0, 1]$. The the number of modes required and the value of the computational cost are normalised with respect to the number of modes and the computational cost for the model at the same Reynolds number and resolving 99% of the turbulent kinetic energy. These two new quantity are indicated as $\hat{N}$ and $\hat{CC}$, respectively. This allows to plot all the models for different values of the Reynolds numbers and resolving different amount of kinetic energy on the same plane. For clarity reasons, different Reynolds numbers are displayed with different symbols. In this analysis we are not interested in the absolute value of the computational cost but in how it scales with the number of degrees of freedom considered in the system after the sparsification is performed. The continuous black line corresponds to the computational cost of the dense system scaling as $\sim N^3$ while the dashed line is the linear interpolation of the markers. Interestingly, it can be observed that for the sparsified models the cost of the evaluation of the quadratic interactions in (4.6) scales down from $\sim N^3$ to $\sim N^{2.5}$. We argue that this could be due to the structure of the energy interactions in modal space where the the band of the dominant relevant interactions (see contours of figure 5.20) identified by the regression grows slowly with respect to the increment of modes needed to keep the energy resolution constant as the Reynolds number increases.

5.3.2 Temporal behaviour of the sparse systems

We now turn our attention to the behaviour of the sparsified models under temporal integration. We consider results for models resolving 95% of the turbulent kinetic energy as an illustrative example, and use the projections of the POD modes onto one of the DNS snapshots to obtain initial conditions. Results are shown in figure 5.32 showing time histories of the turbulent kinetic energy for the three different Reynolds numbers considered, DNS data, displayed in red, is used as reference. Two different reduced order models are selected from all the possibilities: the optimally sparse model (black continuous line), and the dense model obtained by Galerkin-projection (black dotted line). In addition, for the configuration $Re = 2 \times 10^4$ the time dynamics of the sparse system obtained with the greedy approach is also displayed as a black dashed line. As expected, it is found that dense models overestimate the turbulent kinetic energy predicted by the DNS by several orders of magnitude. The over-prediction occurs because the truncation of the small scales in the expression (4.3) leads to a significant imbalance of the production–dissipation budget within the model, this is a well known issue of POD based Galerkin systems well documented in literature (Balajewicz et al., 2013; Noack et al., 2005, 2008, 2011). A qualitatively similar behaviour, if not worse, is then necessarily observed for the model sparsified with the greedy approach, since neglecting weak interactions alone does not cure the original dissipation problems. Conversely, the l_1 model predicts the correct average fluctuation kinetic energy and has excellent long-term stability properties, despite this being not enforced in the regression procedure (see Fick et al. (2018)). This improvement of the performances is due to the fact that the l_1 procedure performs an automatic ”prune-then-calibrate” procedure, where weak interactions are first pruned and the remaining active coefficients are then tuned in the optimisation involved in (4.22) to match the reference dynamics. This mechanism is totally analogous to the one explained in figure 5.4 for the Kuramoto-Sivashinsky equation. It is evident from these results that this second step is key to obtain accurate long-term behaviour. It is worth to be pointed out that:

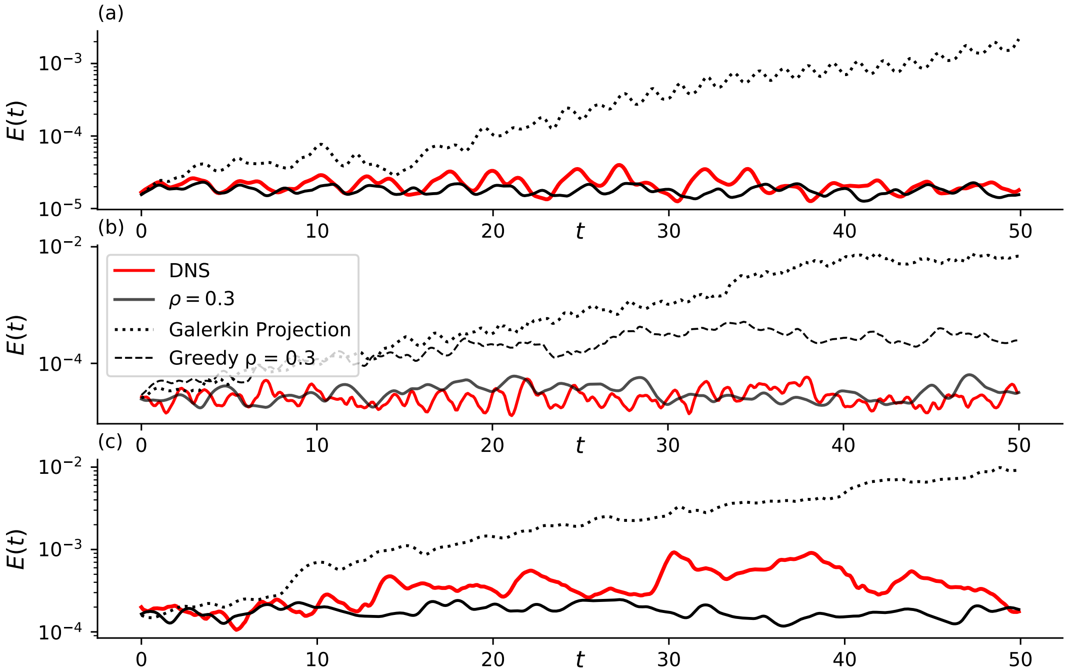

FIGURE 5.32: Instantaneous turbulent kinetic energy $E(t)$ for reduced order models resolving 95% of the turbulent kinetic energy against the DNS (in red). Panel (a),(b) and (c) correspond to $Re = 10^4$,2×10^4 and 5×10^4, respectively. In addition panel (b) shows the behaviour of model sparsified with the Greedy approach as black dashed line.

the sparse system ($\rho = 0.3$) can reproduce the right level of turbulent kinetic energy if the initial condition is close to the DNS dataset. However, we observed that for initial conditions far from the DNS the model does not settle on the correct energy level. This aspect will be discussed later on in this work when the *a priori* sparsification technique is introduced. To better visualise how this prune and tune mechanism implicit in the l_1 regression affects the modal energies of the Galerkin model we visualise the trajectory in phase space and the mean quadratic oscillations, defined as $\lambda_i = \overline{a_i a_i}$ for the models previously considered. In figure 5.33 panels (a) to (c) a shorter segment of state space trajectory projected onto the subspace spanned by the first pairs of modes a_1 and a_2 is shown. Similarly to figure 5.32 the DNS data, the Galerkin model and the l_1 sparsified model are displayed as red, dotted grey and continuous black line, respectively. In addition, for the configuration at $Re = 2 \times 10^4$, we show also results from the temporal integration of the model sparsified with the greedy approach, displayed as blue line in panels (b) and (e). From figure 5.33 panels (a,b,c) a general trend for all the models considered emerges. Namely, the trajectories of the l_1 model remain in the same volume of state space occupied by the DNS projections. This is not true for the projections models drifting away to a different region of state-space, over-predicting the fluctuations amplitude and consequently the integral kinetic energy. This behaviour can be visualised displaying the average modal energy $\lambda_i = \overline{a_i a_i}$ as a function of the modal index in panels (d,e,f). The averages have been computed from $t \geq 40$, corresponding

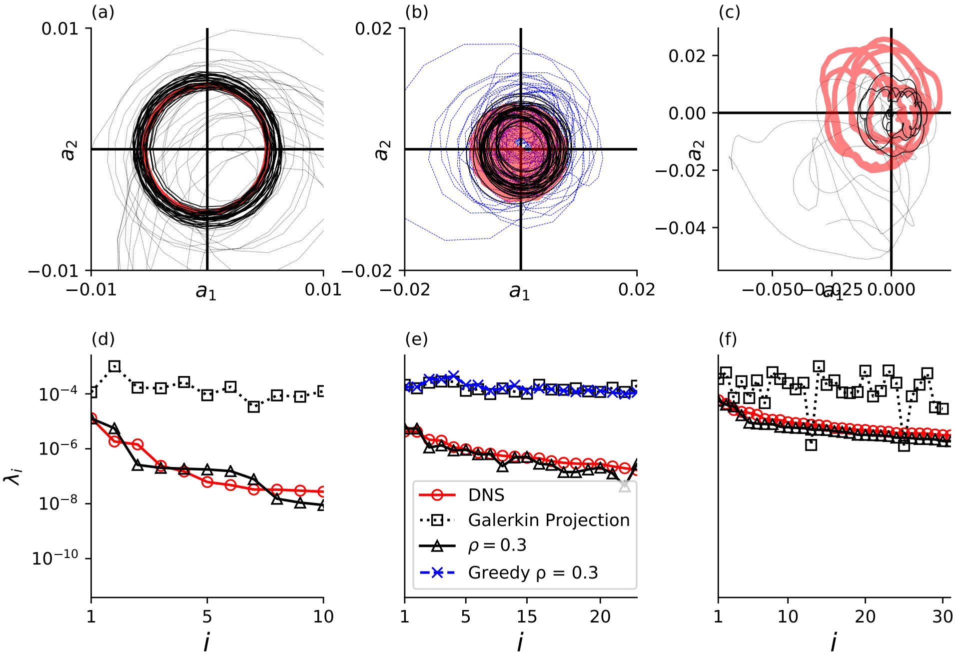

FIGURE 5.33: Orbits in phase space projected on the a_1 ad a_2 plane for the three Reynolds numbers considered in panels (a),(b) and (c), respectively. Panel (d),(e) and (f) show the mean modal amplitude $\lambda_i = \overline{a_i a_i}$. For each Reynolds number the DNS (red), the dense system (grey) and the sparse system (black) are displayed. In addition, for $Re = 2 \times 10^4$ (panels (b) and (e)), the model sparsified with the greedy approach is shown in blue.

to the time needed by the solution to reach statistical convergence as can be observed in figure 5.32. Similarly to the previous analyses DNS data, dense Galerkin model and sparse system are displayed as red circles, black dashed squares, and black triangles, respectively. In addition, for the intermediate value of the Reynolds number results obtained from the integration of the greedy sparsified model is also shown with the blue symbols. Globally it can be observed that the projection and greedy models predict, as expected, much larger energies across the entire spectrum. In addition, almost no or only little decay of the energy is predicted by these models confirming that the truncation of scales implicit in the Galerkin projection prevents the models from predicting the correct viscous dissipation mechanism. Conversely, the l_1 sparsified models correctly predicts the correct levels and decay of modal energies.

To visualise how the sparsification affects the predicted average kinetic energy on all the systems considered, the ratio k_r between the average reference kinetic energy and that predicted by the model is considered. Results are reported in figure 5.34. Each panel displays data for six models, three dense models obtained by Galerkin projection (empty triangles), and the corresponding sparse models with $\rho = \rho_{opt}$ (full black circles). The Reynolds number increases moving from panel (a) to panel (c). A common trend for the three different flow conditions is observed. The models with low energy

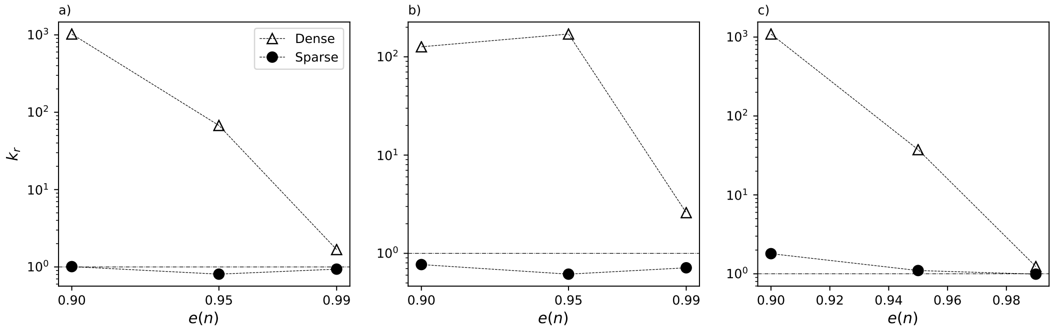

FIGURE 5.34: The ratio k_r for models with different energy resolutions $e(n)$. The empty symbols represent dense systems while the full symbols represent the optimally sparse systems. Results for different Reynolds numbers $Re = 10^4, 2 \times 10^4$ and 5×10^4 are shown in panels a), b), and c), respectively

resolution overestimate the fluctuating kinetic energy up to two orders of magnitude. As the energy resolution of the model increases the value of k_r tends to values close to 1. This result is expected since the number of spatial structures included in the model increases allowing the model to better describe the energy dissipation scales. Interestingly, for all sparsified models, the ratio k_r is close to 1. This means that on average the sparse system predict the right average amount of kinetic energy. In addition, we observe that the ratio k_r is almost constant as the energy resolution increases. This means that regardless the energy resolution chosen the optimisation problem (4.22) is able to provide correct predictions on the kinetic energy of the system. This is a direct consequence to the fact that the target of the sparse regression is the modal accelerations $\dot{a}_i(t)$ computed from DNS simulation. These results shown that it is possible to efficiently use the l_1 sparsification framework to tune the temporal predictions of POD Galerkin reduced order models. The methodology was applied on reduced order models of increasing physical and mathematical complexity, obtaining good predictions on the average level and amplitude of the fluctuations of the turbulent kinetic energy at least for time spans comparable with the training dataset. Crucially, it is, of course, not possible to guarantee that l_1 sparsified models of generic turbulent flows will have good long-term stability (Schlegel and Noack, 2015), but the present results constitute evidence that this is realistically possible on a non-trivial problem.

5.3.3 Reconstructed Flow Field

Once a solution of (4.22) is obtained, rearranging the coefficients vector $\boldsymbol{\beta}_i$ into a new set of matrices $\mathbf{C}_i^s$,$\mathbf{L}^s$ and $\mathbf{Q}^s$ the sparse representation of (4.6) is obtained. The new dynamical system can be integrated to obtain the evolution of a new set of $a_i(t)$ that are used to reconstruct the flow field according to (4.3) since the spatial modes remain unchanged. Figure 5.35 shows the reconstructed vorticity field $\boldsymbol{\omega}$ for the optimally sparse

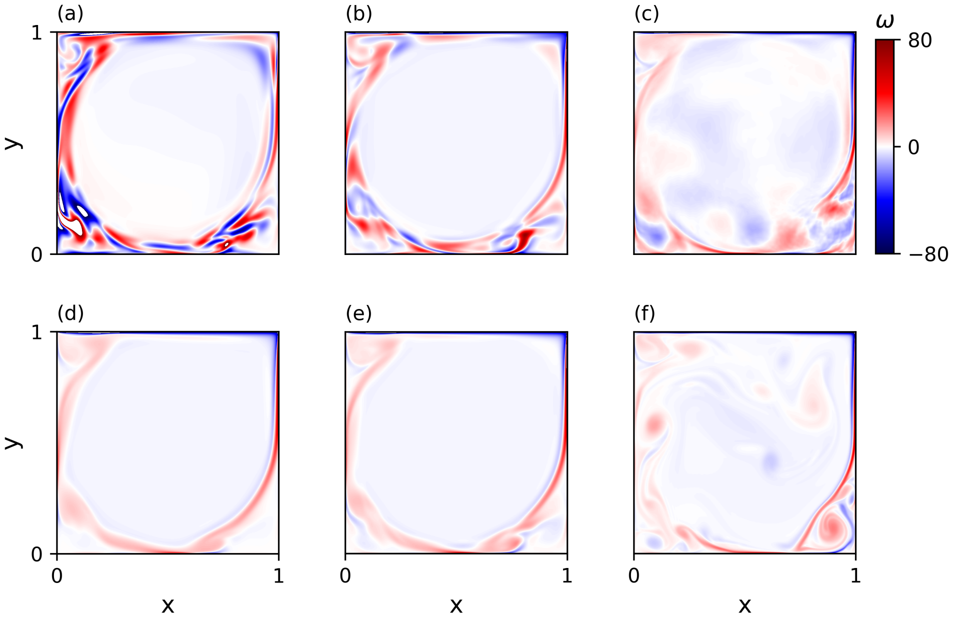

FIGURE 5.35: Instantaneous snapshots of the vorticity field ω. Each column reports data for a different Reynolds number (1×10^4, 2×10^4 and 5×10^4, from left to right). The upper row reports results for the dense reduced order model, while the lower row for the optimally sparse reduced order model.

models with $e(n) = 0.95$. Every column contains results for the three considered values of the Reynolds number, increasing from left to right. The top row shows result of the time integration of the dense system while the bottom row shows the ones of the corresponding sparse model. From a visual analysis, we observe that the overshoot of kinetic energy observed in the dense models in figure 5.32 corresponds to non-physical large amplitude oscillations in the shear layer, as shown in panels (a,b), similarly for the model at highest Reynolds number the vortex in the lower right corner is amplified as well as shown in panel (c). This behaviour is due to an overestimation of the amplitude corresponding to modal fluctuation all over the spectrum. Conversely, sparse models reconstruct quite accurately the topology of the original flow field reproducing the shear layer and the vortex detachment for the lowest Reynolds numbers considered in panel (d) and (e). Similarly, for the highest Reynolds number, panel (f), the sparse model describes the formation of the recirculation vortex in the bottom right corner of the cavity and its shedding and advection along the shear layer by the mean flow. A more in depth analysis on the dynamics of the velocity fluctuations can be done analysing the time average off diagonal Reynolds stress $\overline{u'v'}$. This term is particular relevant for our analysis since the sparsification is performed on the set of equations (4.6) describing only the fluctuating part of the velocity field, while the mean field is left unchanged. Figure 5.36 shows the profile of the Reynolds stress, $\overline{u'v'}$, along a vertical line passing through the centre of the cavity ($x = 0.5$) for increasing values of the

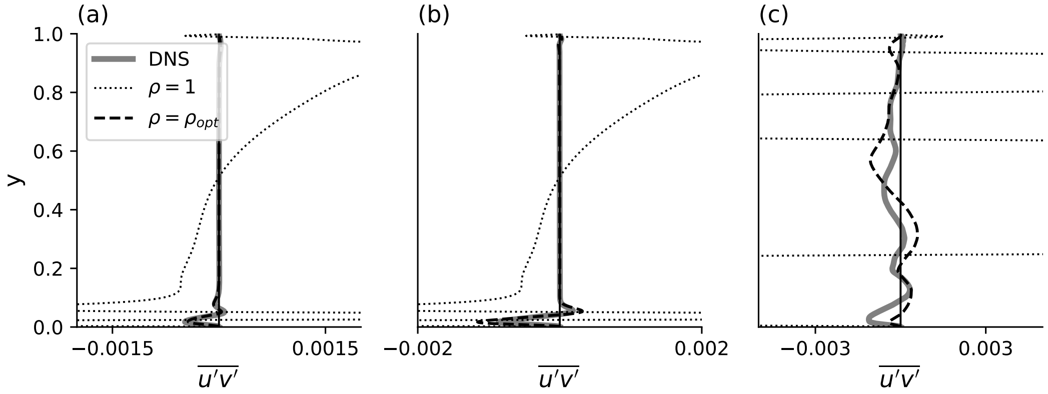

FIGURE 5.36: Profiles of the Reynolds stress $\overline{u'v'}$ from simulation and from time integration of the sparse and dense ROMs along the line $x = 0.5$: in (a), for $Re = 10^4$; in (b), $Re = 2 \times 10^4$; and in (c), $Re = 5 \times 10^4$.

Reynolds numbers, from panel (a) to (c). The grey continuous line is the reference from DNS while the optimally sparse model is represented as black dashed line. Results of the dense ROMs are reported as black dotted line in the background as reference. Since they are orders of magnitude larger than the ones observed in the DNS we chose a range of values on the x axis in order to better visualise the results of the DNS and the sparse ROM.

For the two lower Reynolds numbers, panels (a,b), oscillations in the the stress $\overline{u'v'}$ can be observed, becoming more intense as the Reynolds number increases. In addition, outside the shear layer the stress $\overline{u'v'}$ drops to zero except near the cavity lid. This result is expected since from a visual examination of figure 5.9 and 5.35 we observe that outside the shear layer the flow is mainly stationary and presents quasi-periodic fluctuations. A different behaviour is shown in panel (c) for the highest Reynolds number. As already observed in figure 5.9 and 5.35, in this case the shedding of the bottom-right vortex generates a much more complex flow field, with oscillations extended also near the central area of the cavity. For this flow configuration we observe that although the general distribution is reproduced the values are locally more different with respect to the other two test cases. We argue that this is consequence to the more complex structure of the flow. In fact, since the regression evaluates which terms are statistically relevant the overall behaviour of the fluctuations are reproduced but the spatial distribution can be slightly missed by the sparse system.

5.4 *A posteriori-$l1$* sparsification/calibration of a 18° AoA NACA 0012

In the present section we apply the l_1 sparsification framework to the external flow around a NACA 0012 at high angle of attack. The aim of this chapter is to use the l_1

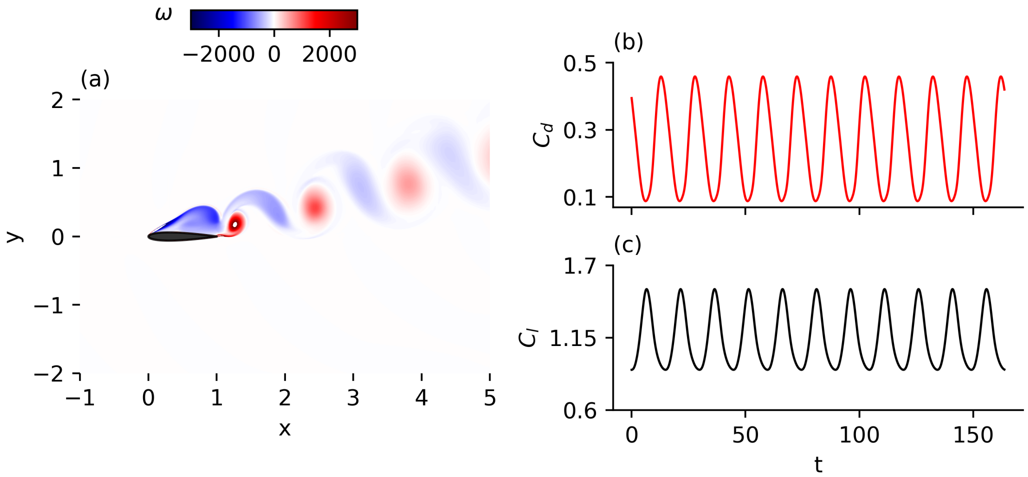

FIGURE 5.37: Panel-(a): vortex shedding around the profile shown by the vorticity field. Time evolution of the drag and lift coefficients in panels (b) and (c), respectively.

framework to tune the coefficients of the reduced order model and test the capabilities of the sparse reduced order model in predicting both local flow parameters, such as vorticity and pressure field, and global parameters such as the aerodynamics coefficients. In addition, due the non trivial geometry it was necessary to develop an interface between the CFD code and the reduced order modelling framework to extract from the mesh the information needed to compute the inner products (4.1) to obtain the matrix coefficients of (4.6). The computational domain has a distance between the far-field boundary and the airflow roughly 8 times the length of the chord to make the results independent from the boundaries. In addition, the grid has been refined in the proximity of the profile with particular attention to the leading and trailing edge areas. The grid independence has been reached on a multi-blocks structured mesh composed by roughly 10^6 elements. An a posteriori check on the solution showed that that the $y+$ is smaller than one for every cell in the boundary layer to avoid the introduction of wall function. The flow conditions have been chosen just below the compressibility limit with Reynolds and Mach number equal to, $Re = 10^5$ and $Ma = 0.25$, respectively. Although, not being an easy scenario for applying our methodology, these flow conditions have been chosen to validate the simulations with a test case already analysed in the research group. In addition, this is a good configuration to test the calibration framework for flow conditions at the boundaries of the applicability of the reduced order modelling techniques explained in the present work.

Due to the high Reynolds number for this configurations a URANS simulation with the solver `pimpleFOAM` has been performed. The k-ω turbulence model has been used. The solver has been modified to save to disk the Eulerian acceleration field at each saved timestep. Figure 5.37-(a) shows a snapshots of the out of plane vorticity field ω_z around the NACA airfoil. The high angle of attach causes an early stall on the suction side of the airfoil, leading to flow separation producing unsteady vortex shedding advected

n	1	2	3	5	10	20	30	40	50
POD	0.48	0.97	0.976	0.98	0.9	0.95	0.98	0.995	0.998

TABLE 5.3: Normalised cumulative energy distribution $e(n)$ for the POD decomposition.

by the main flow as shown in panel (a). The vortex shedding generates oscillation of the pressure field on the surface of the profile resulting in oscillating aerodynamic forces. Panel (b) and (c) of figure 5.37 show the temporal evolution of the aerodynamic coefficients C_d and C_l defined as the integral over the profile of the pressure and shear field in the direction parallel and orthogonal with respect to the freestream. In this work the aerodynamic coefficients are computed using the OpenFOAM standard libraries for post-processing *libforces.so*.

5.4.1 Modal decomposition

To generate the Galerkin model we performed the POD decomposition with a POD tool directly implemented in the OpenFOAM framework. Table 5.3 shows the cumulative value of POD eigenvalues λ_i. It can be observed that the first eigenvalues recovers already 48% of the fluctuating kinetic energy and the first pairs recovers almost 98%. Note that the first two eigenvalues have the same energy content, meaning that the first two POD modes describe an oscillating structure. This flow structure is the unsteady vortex shedding downstream of the airfoil. In addition, we observe that only little energy is contained in the remaining modes. This is consequence of the conjoint effect of the strong periodic nature of the flow and the URANS simulation that smears away the smallest flow features. Since information about the smallest and dissipative scales of motion is lost the generation of reliable temporally stable reduced order system is challenging for this configuration. Therefore, the presence of a calibration step absolutely necessary for the stabilisation of the model itself. This is better shown in figure 5.38 displaying the x component u of the velocity field for the first POD mode and its power spectrum in panel (a) and (b), respectively. The same quantities for the 10-th POD mode are shows in panels (c) and (d). As expected the most energetic mode is associated to the large scale structures responsible of the vortex shedding phenomenon and of the detachment of the separation bubble on the suction side of the profile. In addition, we expect the second temporal mode to have the exact same frequency content and describe the same flow structure up to a minus sign. This is due to the strong periodic nature of the flow forcing the POD to produce a pair of modes with the same energy content representing the real and complex part of a complex wave. A similar behaviour was already observed for the POD decomposition of the cavity at the lowest Reynolds numbers. On the other hand, the 10-th mode, shown in panel (c), describes much smaller flow structures associated with the interactions between the freestream

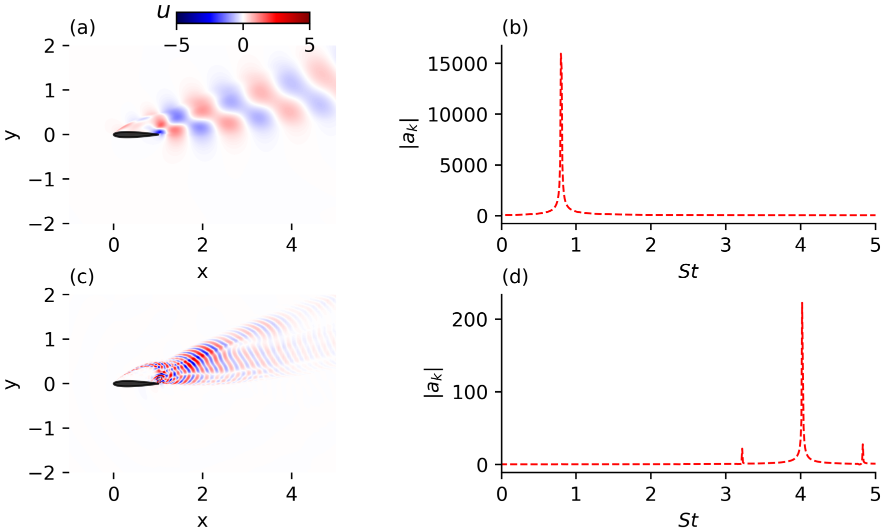

FIGURE 5.38: Horizontal component of the velocity field u for the first and tenth POD mode in panels (a) and (c), respectively. Panels (b) and (d) show the amplitude of the Fourier transform of the first and tenth temporal mode.

and the detachment bubble on the suction side of the profile. These small structures are advected downstream by the larger ones. Therefore, we expect a pattern of triadic interactions (4.16) qualitatively similar to the one observed for the cavity flow. In addition, a connection between the frequency of oscillation of the first POD modes and the oscillation of the lift and drag coefficients can be observed. This supports the idea of the link between the most energetic flow structure describing the vortex shedding and the oscillation of pressure around the profile itself. The reduced order model was generated using the first 10 POD modes. A system formally identical to (4.6) is obtained. It is important to underline that, unlike the cavity flow, in this case the pressure term is different from zero. Following the idea proposed by Cordier et al. (2010) we chose to not include this term in the model and let the sparsification procedure to tune the model coefficients to include the effect of this term originally neglected. This approach is, in a sense, similar to the one proposed in Galletti et al. (2004), where the pressure term is modelled as quadratic combination of the temporal modal coefficients $a_i(t)$.

5.4.2 *A posteriori* tuning and sparsification

In the present case we consider a model composed by $N = 10$ modes resolving more than 99% of kinetic energy as shown in table 5.3. Problem (4.22) is solved for increasing values of the regularisation weight γ_i obtaining a family of sparse models with different accuracy/sparsity trade-off. This work employs the K–fold validation procedure available in the *sklearn* library (Pedregosa and Varoquaux, 2011). However due to the large number of snapshots available with respect to the size of the database $\mathbf{\Theta}(\mathbf{A})$ the standard deviation is small and relatively constant along the sparsification curve $\rho - \epsilon$.

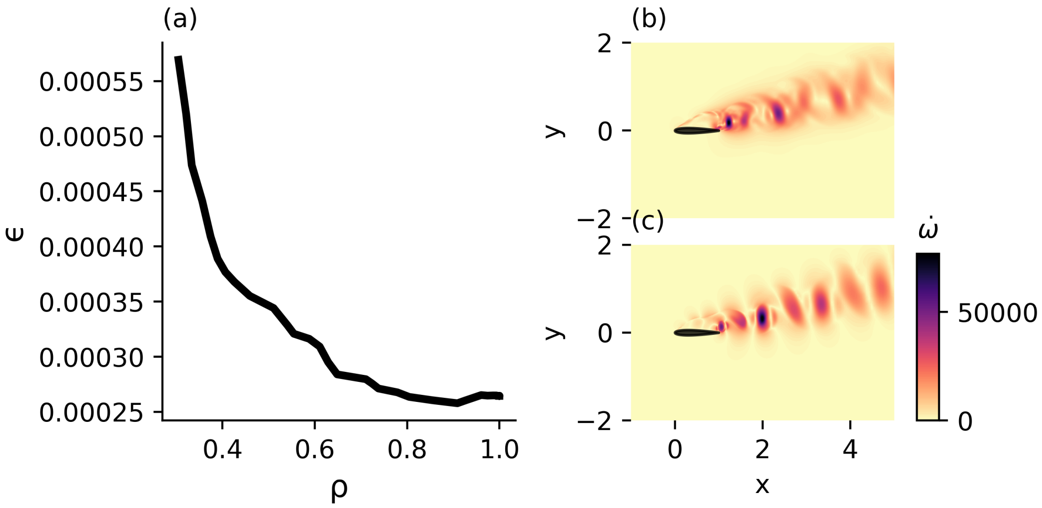

FIGURE 5.39: Panel (a): $\rho - \epsilon$ curve obtained with the *a posteriori* sparsification approach. Absolute value of the Eulerian acceleration field from CFD and for $\rho = 0.6$ in panel (b) and (c), respectively.

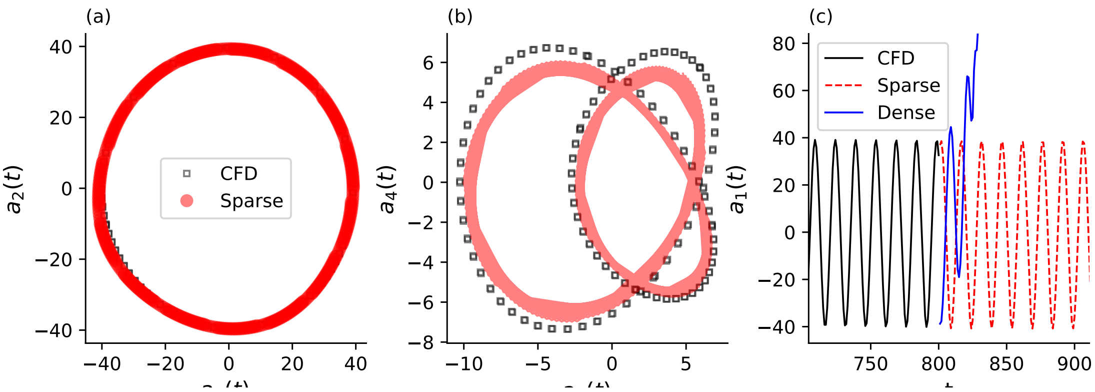

FIGURE 5.40: Time evolution in modal space of the first and second couple of modal coefficient against the CFD in panel (a) and (b), respectively. Panel (c): shows the time evolution of $a_1(t)$ for the sparse the dense model against the CFD used for the training.

Therefore, it has been omitted from figure 5.39-(a), showing the $\rho - \epsilon$ curves for the model considered. For this configuration it can be observed the presence of a small plateau for values of the density in between 0.7 and 1. This is expected and due to the relatively small size of the model considered. In the continuation of the analysis we consider a model with $\rho \sim 0.7$ where the reconstruction error is still low and only the non relevant terms have been removed from the system. Panel (b) and (c) shows the Eulerian acceleration field from CFD and the one reconstructed by the sparse system, respectively. Once the new set of coefficients is obtained by solving problem (4.22) the new dynamical system can be temporal integrated. In the present work we have integrated the system for $T = 200$ time units from $T = 700$. Figure 5.40 panels (a) and (b) show the orbits in phase space for the first two pairs of modes ($a_1 - a_2$ and

$a_3 - a_4$), respectively. We observe that the sparsification procedure tunes the modal amplitude on the right value up to a small decrease of a_3 and a_4 due to the effect of the l_1 regularisation that pushes inevitably the numerical values of the coefficients towards the axes, decreasing their absolute value. Panel (c) shows the temporal evolution of $a_1(t)$ for three different models taken into consideration. Namely, the CFD results, the sparse system and the dense Galerkin system are displayed as black,red dashed and blue lines, respectively. Operatively, this graph has been obtained using the data from CFD to generate the modal decomposition and train the l_1 regression. Once the sparsification is terminated the sparse system can be used to advance the system in time at a much lower computational cost. As expected the dense system quickly overshoots the correct level and subsequently ends up in numerical divergence. Conversely, the sparse system keeps the right oscillation amplitude and frequency. This confirms that the sparsification/calibration step is necessary to obtain reduced order models able to provide reliable long term prediction.

5.4.3 Flow reconstruction

Once the new set of temporal modes are obtained from temporal integration the flow field can be reconstructed as (4.3). Figure 5.41 shows the vorticity fields for two different snapshots for the dense system panels (a-b) and for the sparse system panels (c-d). As expected the dense system quickly diverges leading to nonphysical fluctuations in the wake of the profile. Similarly to what observed previously, this behaviour is direct consequence of inability of the dense system of providing the right amount of dissipation in the less energetic and smaller scales of the model. Consequently, the small structures, such as the modes described in figure 5.38 panels (d) and (e), are over-amplified resulting into an un-physically wiggling flow field leading subsequently to divergence. This behaviour is also aggravated by the fact the dataset is obtained from URANS simulation including no or little information about the small scales of the flow. Conversely, the sparse system generates a stable solution that leads to a flow field coherent with the one obtained by CFD. In panels (c) and (d) it can be observed that the reconstructed solution reproduces faithfully the vortex shedding observed in the CFD simulation showing much better stability properties for the time interval analysed.

5.4.4 Pressure field reconstruction

Reduced order modelling produces a system of ODEs (4.6) describing the temporal evolution of $a_i(t)$. Thus, the velocity field can be reconstructed according to the modal expansion (4.3). However, due to the solenoidal nature of the spatial basis functions $\boldsymbol{\phi}_i(x)$ the incompressibility constraint in (4.2) is automatically satisfied and no equation describing the pressure field is present. Since, for some applications the knowledge

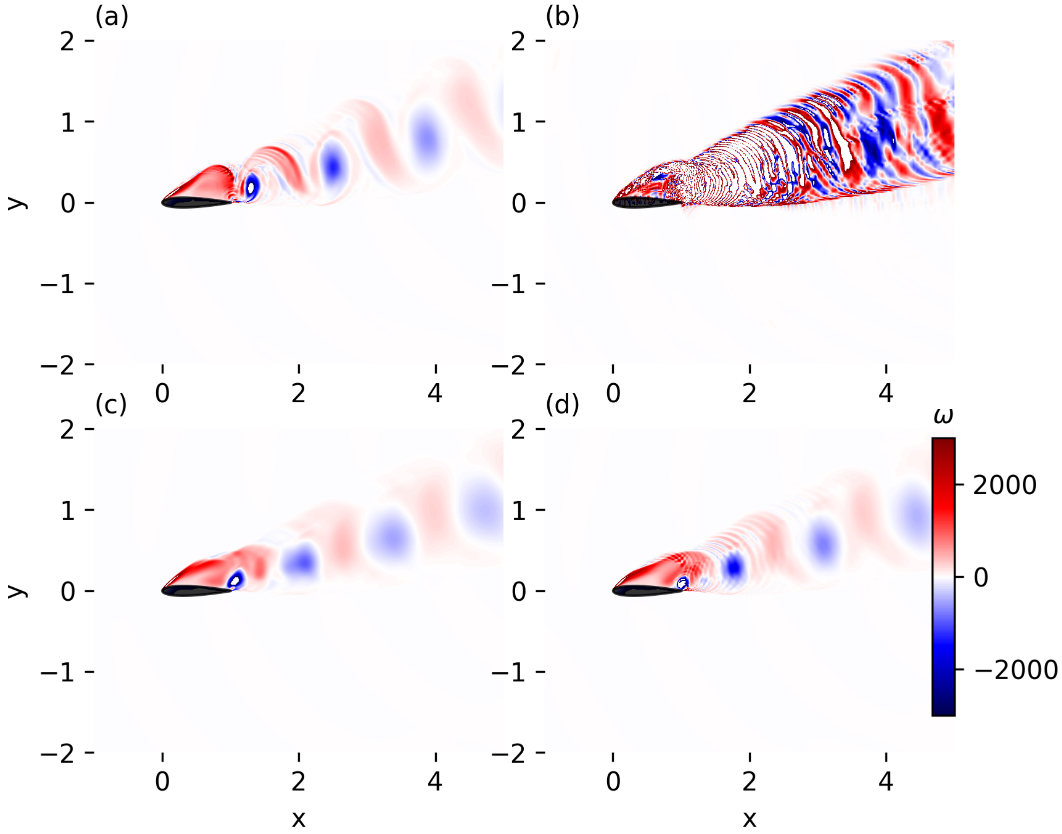

FIGURE 5.41: Snapshots of the vorticity field for $t = 800$ and $t = 850$. Reconstructed flow field for the dense system in panel (a) and (b). Reconstructed flow field for the sparse system in panel (c) and (d).

of the pressure field is of paramount importance several approaches to reconstruct the pressure field a posteriori have been developed (Caiazzo et al., 2014; Noack et al., 2011).

Here, we follow Caiazzo et al. (2014) where the pressure field is reconstructed *a posteriori* using the sparse ROM velocity solution. The idea is to reconstruct the pressure field utilising the pressure Poisson equation obtained taking the divergence of the momentum equation (4.2)

$$\nabla^2 p = \nabla \cdot (\mathbf{u} \cdot \nabla \mathbf{u}) \tag{5.21}$$

The velocity field at the right hand side is now approximated by the ROM velocity field $\mathbf{u}_r$ expressed by the finite series (4.3). Since the basis functions are divergence free equation (5.21) becomes:

$$\nabla^2 p_r = \sum_{i=0}^{N} \sum_{j=0}^{N} a_i(t) a_j(t) (\sum_{n=1}^{d} \sum_{m=1}^{d} \partial_n \boldsymbol{\phi}_i \partial_m \boldsymbol{\phi}_j) \tag{5.22}$$

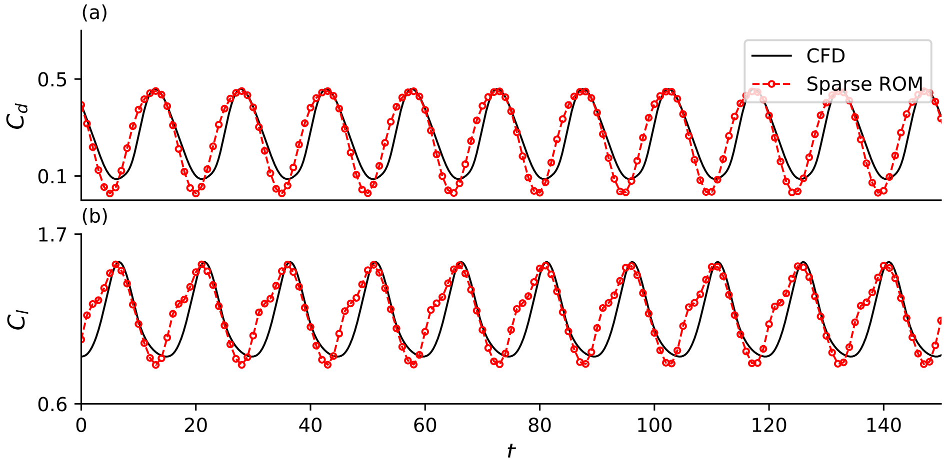

FIGURE 5.42: Temporal behaviour of C_d and C_l in panels (a) and (b), respectively.

where d is the dimension of the physical space we are working in ($d = 2$ in this case). The solution of (5.22) has the form

$$p_r(t, \mathbf{x}) = p_{r0}(\mathbf{x}) \sum_{i=0}^{N} \sum_{j=0}^{N} a_i(t) a_j(t) \tag{5.23}$$

with p_{r0} is obtained solving the following equation

$$\nabla^2 p_{r0} = \sum_{n=1}^{d} \sum_{m=1}^{d} \partial_n \boldsymbol{\phi}_i(\mathbf{x}) \partial_m \boldsymbol{\phi}_j(\mathbf{x}) \tag{5.24}$$

Since the spatial basis functions are not time dependent p_{r0} can be computed in a pre-processing step and then utilised for the solution of (5.23). Subsequently, p_r can be computed at each time step once the temporal coefficients $a_i(t)$ are available. In the present case, the reconstructed pressure field is reconstructed at each time step of the velocity ROM. The pressure and velocity fields are then utilised to compute the aerodynamics coefficients C_d and C_l through the OpenFOAM standard libraries on the reconstructed flow field. Figure 5.42 shows the temporal behaviour of the drag and lift coefficients C_d and C_l in panel (a) and (b), respectively. The black line corresponds to the CFD predictions while the red symbols are the values predicted by the sparse ROM. Since, after the pre-processing step, only the instantaneous value of $a_i(t)$ is needed the spacing of the reconstructed C_d and C_l can be much wider than to the one obtained from the CFD. In can be observed that the sparse system reproduces the correct frequency of oscillations of both C_d and C_l. In addition, in panel (a) we observe that the maximum value of C_d is predicted quite accurately while the minimum is slightly underestimated leading to a small over-prediction of the fluctuations amplitude. Similarly, in panel (b) we observe a good accuracy of both the maximum and minimum value of C_l. We observe a periodic discrepancy in the rising part of the C_l curve, we argue that this could be due to the relatively small number of modes chosen to generate the model.

5.5 Summary: *a posteriori* sparsification

In this chapter the *a posteriori* sparsification and tuning of reduced order models have been discussed and tested on models of increasing physical and mathematical complexity. This section aims to briefly collect all the major results and discuss limitations and possible improvements of this approach. First, the one dimensional spatially extended system described by the Kuramoto-Sivashinky equation was considered. The key outcomes of this study are the followings.

- The quadratic interaction tensor **Q**, restricted to each mode, has entries of constant intensity $k/2$. However, the triadic interactions tensor contains a subset of dominant interactions.
- The l_1 framework can successfully identify this subset in accordance with the strength of the energy interactions reconstructing the correct non locality of triadic interactions observed in the energy analysis.
- Two different approaches for the choice of the regularisation weight γ_k have been proposed. One consists in keeping the regularisation weight γ_k constant across the modes, the other varies the intensity of the regularisation weight γ_k according to the amplitude of the mode itself. Results show that these two choices identify a pattern of interactions in agreement with the absolute and the relative strength of the energetic interactions, respectively. In addition, we observed that the first approach generates the lowest reconstruction error for the same density.

Subsequently, the flow developing inside a lid-driven square cavity has been considered. This configuration is particularly convenient since the closed domain simplifies the generation of the model (4.6) getting rid of the pressure term. More importantly, despite the simple geometry, models of increasing complexity can be easily obtained. One of the goals was to understand the role of the subspace utilised for projection on the structure and sparsity of energy interactions between modes. The key findings of this analysis are the followings.

- The analysis of the average energy transfer rates between modal structures has shown that, for both the POD and DFT-based models, a subset of most relevant interactions exists. This is an *a-posteriori* feature of solutions of the equations and not an *a-priori* property of the evolution equations. The model coefficients identified by the Galerkin projection do not have a particular structure and are typically different from zero.
- In both cases, there exists a sweet-spot on the $\rho - \epsilon$ curve where the sparsification approach recovers correctly this subset, with little effect on the prediction accuracy. The models also preserve to a good degree of accuracy the non-local nature

of triadic interactions and the conservation properties of the convective term of the Navier-Stokes equations.

- The effectiveness of the sparsification grows both with the number of modes utilised in the projection (energy resolution) and with the Reynolds number of the flow considered. The interesting consequence is that, while the total number of quadratic interactions grows cubically with the number of modes, the number of relevant interactions does not grow as quickly as a consequence of the structure of energy interactions in modal space.
- The sparsity pattern is not invariant under different modal representations of the flow. Reduced-order models obtained with DFT modes show a different structure of the triadic interactions tensor with a much higher sparsity. However, the sparsification framework can identify the correct non local physics of the interactions, regardless the modal representation chosen. In addition, this feature suggests that it might be possible to mathematically define a modal basis that maximises the sparsity of triadic interactions.

Lastly, a test case of more practical interest was considered. Namely, the separated flow around a NACA 0012 profile at a high angle of attack. For this flow configuration, we focus on the ability of the LASSO regression to sparsify and, more importantly, tune the predictions of the ROM.

- Conversely to what was observed for the dense system the sparse/tuned system shows good prediction capabilities under temporal integration. The results are in agreement with the CFD both in time domain and phase space.
- The velocity and the pressure fields have been reconstructed from the temporal coefficients obtained by integration of the sparse/tuned model. It can be observed that the main features of the flow are reproduced.
- The behaviour of the aerodynamics coefficients in the reconstructed flow is reproduced accurately both from the point of view of the amplitude and oscillation frequency

Chapter 6

Results: *a priori* sparsification

The *a posteriori* sparsification involves a ”prune and tune” procedure applied on the system coefficients by the LASSO regression. As shown in the previous chapter this approach could lead to some difficulties in the interpretation of the role of flow structures in the energy transfer, since triadic interactions are modified while the spatial basis functions are kept unchanged. Starting from the key observation that the sparsity pattern is not invariant with respect to the modal basis chosen to perform Galerkin projection, a novel sparsification methodology, called *a priori* sparsification, encoding the sparsity features of the flow in the modal basis, has been developed. In this chapter, this methodology is applied to the chaotic two dimensional lid-driven cavity flow. This configuration, already analysed in the previous chapter, allows a comparison between the two different sparsification methodologies.

6.1 Two dimensional lid-driven cavity flow

In this section, the *a priori* sparsification framework has been applied to the two dimensional lid-driven cavity flow at $Re = 2 \times 10^4$. Before moving on to the results of the sparsification, in the following sections, some preliminary analyses about the sparsity promoting effects of the l_1 norm and the effect of the initial guess on the solution are shown. The results shown in this section are currently under consideration for a publication Rubini et al. (2021).

6.1.1 Sparsity-promoting effect of the l_1-based constraint

To understand the role of the l_1 norm in the solution of (4.35), a formal discussion based on the proximity operator theory can be formulated, using the classical LASSO formulation as a starting point (Friedman et al., 2008). Here, to understand heuristically

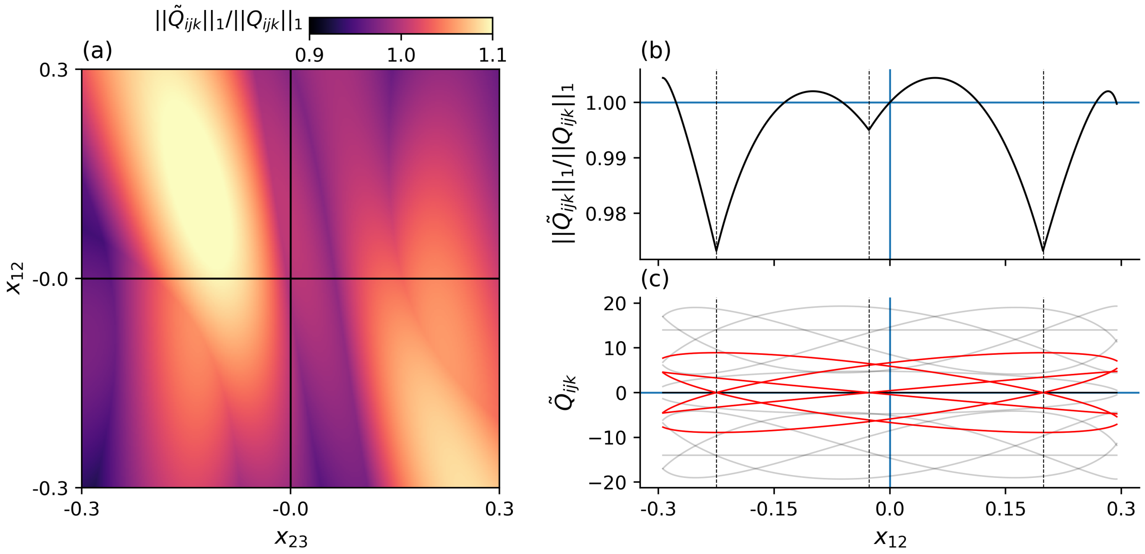

FIGURE 6.1: Map of $||\tilde{Q}_{ijk}||_1/||Q_{ijk}||_1$ projected on the $X_{12} - X_{23}$, panel (a). One dimensional cut of the same quantity by varying the coordinate X_{12}, panel (b). The twenty-seven entries $\tilde{Q}_{ijk}$ as a function of X_{12}. The six entries that vanish identically in correspondence of the spikes of $||\tilde{Q}_{ijk}||_1/||Q_{ijk}||_1$ are highlighted in red.

the role of constraint, as an example, we consider a small reduced-order model of the lid-driven cavity flow at $Re = 2 \times 10^4$. To easily visualise the effect of the rotation on the properties of the coefficient tensor $\mathbf{Q}$, we consider a small model with dimensions $M = 5$ and $N = 3$ and perform a parametric study on the effects of some of the entries of $\mathbf{X}$ on the quantity $||\tilde{\mathbf{Q}}||_1$ without solving (4.35), but simply applying the definition (4.27). Figure 6.1-(a) shows contours of $||\tilde{\mathbf{Q}}||_1/||\mathbf{Q}||_1$ on the $X_{12} - X_{23}$ plane, with all other off-diagonal entries of $\mathbf{X}$ set to zero. Several sharp-bottomed valleys and local minima can be observed, arising from the non-convex, non-smooth nature of the quantity $||\tilde{\mathbf{Q}}||_1$. Panel (b) shows a cut of $||\tilde{\mathbf{Q}}||_1/||\mathbf{Q}||_1$, shown in panel (a), along the x_{12} coordinate. Here too, we observe a similar structure where the minima are identified in the plot by the dashed vertical lines. More interestingly, each minimum corresponds to a point with discontinuous first derivative, due to the nature of the l_1 norm and some entries $\tilde{Q}_{ijk}$ crossing the zero axis. This behaviour is illustrated in panel (c) where the entries of $\tilde{\mathbf{Q}}$ are shown as a function of X_{12}, with those crossing the zero for some value of X_{12} highlighted in red. The general idea is that increasing the penalisation weight ξ in the sparsity promoting constraint (4.35)-(c) pushes $||\tilde{\mathbf{Q}}||_1$ closer and closer to one these valleys. In practice, some entries $\tilde{Q}_{ijk}$ are forced to zero due to the strong gradient at the point of non-differentiability, resulting in a sparse coefficient tensor $\tilde{\mathbf{Q}}$. However, unlike in the *a-posteriori* LASSO-based sparsification method, not all quadratic coefficients can be simultaneously set to zero by an arbitrary rotation, as is clear from panel (c). In fact, the tensor $\tilde{\mathbf{Q}}$ depends nonlinearly on the rotation $\mathbf{X}$, while in the LASSO-based method the tensor coefficients are directly the optimisation variables of the problem.

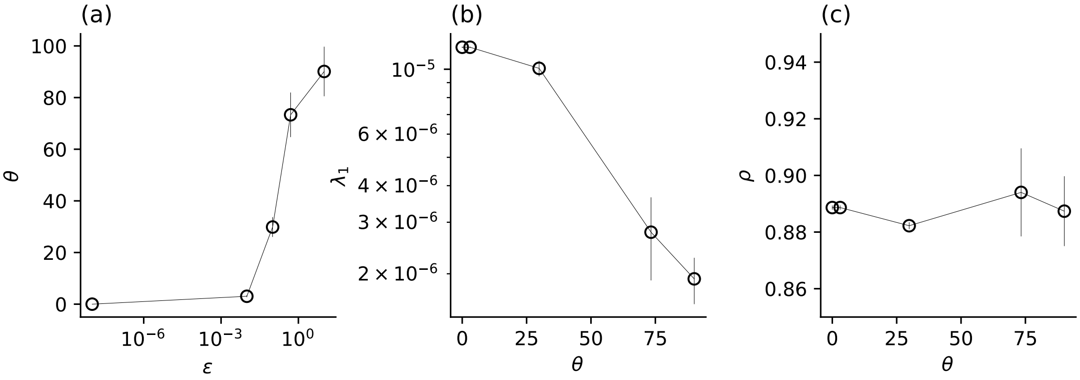

FIGURE 6.2: Panel (a): relation between the value of the perturbation ϵ and the average rotation angle between sub-spaces θ . Mean values and standard deviation for the first eigenvalue λ_1 and density ρ for the different system obtained, in panel (b) and (c) respectively.

6.1.2 Dependence of the solution from the initial guess

Problem (4.35) is non-convex due to the constraint (4.35b). This introduces the problem of characterising the uniqueness and the quality of the solution with respect to other local minima. To understand how this problem affects the optimisation results, we considered different sets of initial guesses defined as increasingly larger perturbations of the identity, *i.e.* as:

$$\mathbf{X}_0 = \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \end{bmatrix} + \epsilon \mathbf{R}, \tag{6.1}$$

with $\mathbf{R} \in \Re^{M \times N}$ a randomly generated rotation matrix, satisfying $\mathbf{R}^\top \mathbf{R} = \mathbf{I}$, with $\mathbf{I} \in \Re^{N \times N}$. We generated ten random rotation matrices for different increasing values of ϵ in the range $[10^{-8}, 10]$ and solved problem (4.35) keeping $\xi = 2$ to target an intermediate density and ensure to fall into the feasibility region of the sparsity constraint. Figure 6.1.2-(a) shows the effects of ϵ on the average rotation angle θ between columns of the initial guess matrix $\mathbf{X}_0$ and column of the solution matrix obtained with increasing values of the perturbation ϵ. Circles and vertical bars identify the average and standard deviation of this quantity across the ten different samples. In panels (b) and (c), we show these statistics for the eigenvalue λ_1 and the density ρ, respectively, as a function of the θ. It can be observed that for low values of θ all initial guesses converge to the same solution, as it can be observed from the low values of the standard deviation. Conversely, when initial guesses starting from regions further away from the identity the standard deviation starts growing, as problem 4.35 converges to different solutions. This result is expected and it is in agreement with the shape of the sparsity constraint shown in figure 6.1-(a). Crucially, this result supports the initial hypothesis of looking for solutions as small rotations away from the POD basis. Results shown later on have thus been obtained using this initial hypothesis. This agrees with our original motivation, *i.e.* to obtain a basis as close as possible to the energy optimal

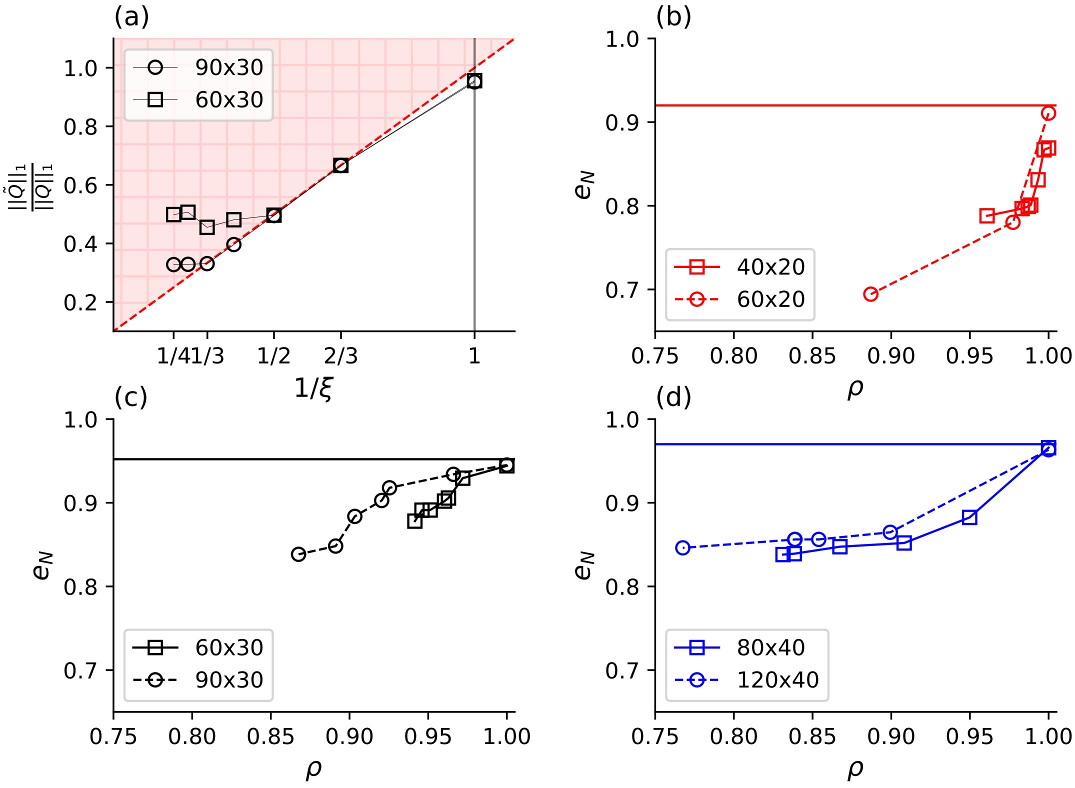

FIGURE 6.3: Panel (a): visualisation of the sparsity-promoting constraint 4.35b on the plane ξ^{-1}-$||\tilde{\mathbf{Q}}||_1/||\mathbf{Q}||_1$. The red region denotes the infeasible set. Panels (b) to (d): $\rho - e_N$ curves for families of models with dimension $N = 30, 20, 40$ and ratios $M/N = 2, 3$ (open squares and open circles, respectively).

POD that leads to a sparser energy interaction pattern. Lastly, we observe that λ_1 decreases monotonically with the size of the perturbation since the farther is the initial guess from the identity matrix, the larger is the loss of energetic optimality. A different behaviour is observed for the density where the average value is comparable for all initial guesses considered.

6.1.3 Model sparsification

We consider families of models with $N = 20, 30$ and 40 (see table 5.1 for the energy resolution), generated from rotations of the original POD basis modes in subspaces of increasing dimension M. Note that when the starting subspace dimension M is equal to N, the stability constraint (4.35b) can never be satisfied, *i.e.* no solution is acceptable or useful. This behaviour was observed by Balajewicz et al. (2016), who pointed out that the case $M = N$ must be discarded since no temporal stabilisation can be achieved. Hence, we show here results for ratios $M/N = 2, 3$. Because of the definition of problem (4.35), a family of models with different densities ρ and reconstructed kinetic energy can be generated by varying the penalisation parameter ξ. In practice, this is achieved by progressively increasing ξ starting from 1, where the sparsity-promoting constraint is not active, and obtaining η iteratively to ensure long-term temporal accuracy. In addition, has been observed that the optimal value of η is dependent on the energy resolution N. This was expected since, for larger systems more scales are considered, consequently more dissipation is introduced in the system and a smaller correction is required. However, the variation of η between different systems can be considered to be small.

We first consider models generated with $N = 30$ basis modes for increasing ratios M/N. Figure 6.3-(a) displays optimal solutions on the plane ξ^{-1}-$||\tilde{\mathbf{Q}}||_1/||\mathbf{Q}||_1$. The red dashed line separates solutions that satisfy the sparsity-promoting constraint (4.35b) (white area) from solutions that do not (red area). Note that all solutions satisfy the stability constraint, so the white region will be referred to as the feasibility region. For the sake of completeness, we report that also solutions for $M/N = 1$ have been computed, however due to the square nature of the rotation matrix $\mathbf{X}$, they do not satisfy the stability constraint and are meaningless for our discussion. Therefore we chose not to display them in figure 6.3. Considering data points for $M/N = 2$, for small ξ, the optimisation problem has feasible solutions generally falling on the boundary of the feasibility region. However, there exists a thresholds value of ξ above which the optimisation problem terminates unsuccessfully in the unfeasible region. This is manifested in panel (a) by a sudden turn of the solution traces from the feasible region boundary upwards into the red region. Crucially, this threshold value increases for larger ratios M/N, *i.e.*, a more pronounced sparsification can be achieved for larger dimensions of the starting subspace. This is arguably a consequence of the fact that higher ratios M/N correspond to more degrees of freedom available in 4.35 to ensure that the sparsity constraint is satisfied. From a mathematical point of view, this behaviour is a direct consequence of the fact that a simple rotation of the basis modes, defined by the matrix $\mathbf{X}$, cannot shrink to zero an arbitrarily large number of coefficients in the tensor $\mathbf{Q}$. To better visualise this property we can plot the trend of the entries of the rotated tensor $\tilde{Q}_{ijk}$ as the regularisation weight ξ is increased. Figure 6.4-(a) is analogous to panel (a) in figure 6.3 and it is reported only for clarity of visualisation. Panel-(b) displays the entries of $\tilde{Q}_{ijk}$ against ξ^{-1}. We observe that while decreasing the values of $||\tilde{\mathbf{Q}}||_1/||\mathbf{Q}||_1$ because of the sparsity promoting constraint the entries of $\tilde{Q}_{ijk}$ are continuously decreased as well. More interestingly, we observe that this behaviour does not continue indefinitely since all the entries become zero. Conversely, at a certain value of the regularisation weight ξ^{-1}, $||\tilde{\mathbf{Q}}||_1/||\mathbf{Q}||_1$ reaches a plateau. Similar behaviour can be observed for the entries $\tilde{Q}_{ijk}$, decreasing to a certain value and then remaining approximately constant. This results in some entries shrunk to zero and others not (shown by the two red lines) and the consequent existence of the minimal value for the density ρ as discussed previously.

At a first glance, the fact that a minimum value of density different from zero exists for any given system may appear as a limitation of the present *a priori* method. However,

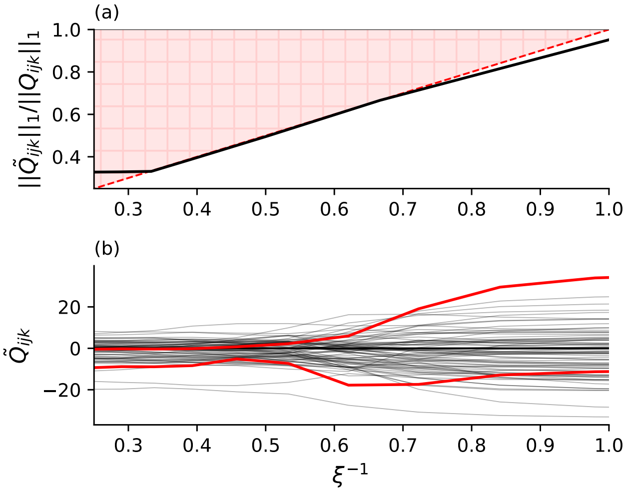

FIGURE 6.4: Evolution of $||\tilde{Q}_{ijk}||_1/||Q_{ijk}||_1$ and entries of the rotated quadratic interactions tensor $\tilde{Q}_{ijk}$ against the regularisation weight ξ in panels (a) and (b), respectively. The regularisation weight increases from right to left.

this is an advantage with respect to the *a posteriori* LASSO-based model sparsification technique. In that case, feasible solutions could always be found for arbitrarily large regularisation weights, where all system coefficients are shrunk to zero in the limiting case. These models, however, were observed to have little physical significance. For instance, in the limiting case with zero density the models would not capture any temporal fluctuation, *i.e.* resulting in the trivial system $\dot{a}_i(t) = 0, i = 1, \dots, N$. With the present *a-priori* sparsification method, all feasible systems are temporally stable and provide physically consistent predictions. Therefore, there is no ambiguity linked to the choice of ξ for a given M/N, since the feasible model obtained at the largest ξ (smaller density) should be selected. In what follows we will adopt this approach.

When the penalisation ξ is increased, sparser models are obtained from the optimisation through small rotations of the original POD basis. This, however, comes at the cost of decreasing the energy optimality of the resulting basis. This effect can be visualised on the $\rho - e_N$ plane, as shown in figures 6.3-(b,c,d), for $N = 20, 30, 40$, respectively. The horizontal line in each panel corresponds to the fraction of reconstructed energy of the original dense POD-Galerkin projection model. In these panels, squares are used to denote data for $M/N = 2$, while circles denote data for $M/N = 3$. First, we observe that systems for $\rho = 1$ ($\xi = 1$) do not reconstruct the entire kinetic energy captured by the original POD basis. This behaviour is because the presence of the stability constraint (4.35b), even in the configuration of no sparsification $\xi = 1$, produces a slight rotation of the new basis. When ξ is increased lower densities and lower reconstructed kinetic energy are obtained. However, the model dimension affects the slopes of the curves, *i.e.* the larger the model size, the more the model can be sparsified without affecting the ability of the new basis to reconstruct kinetic energy. This indicates that the sparsification technique becomes more effective as the model complexity increases.

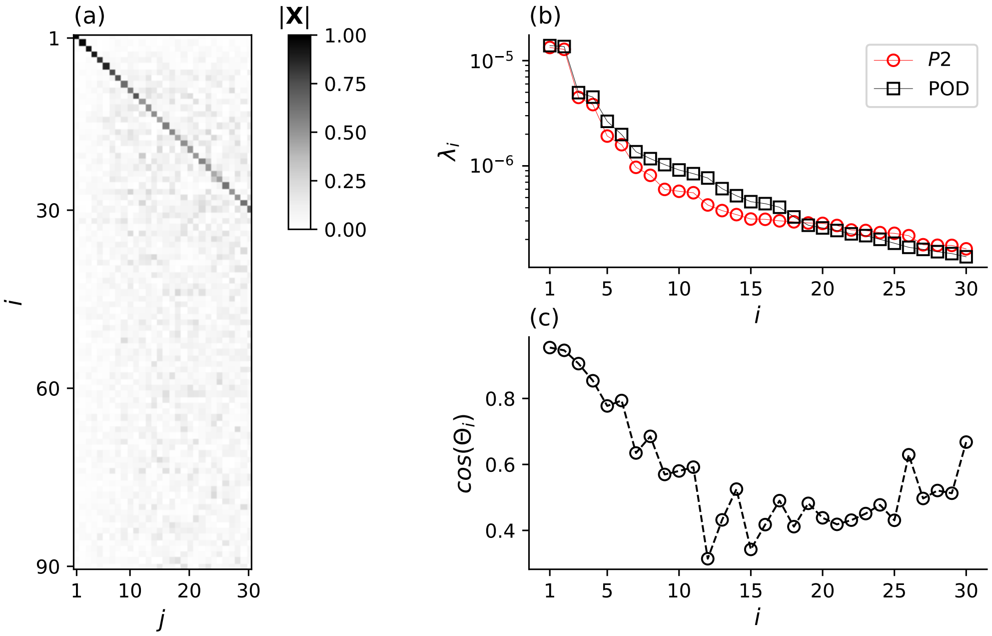

FIGURE 6.5: Optimisation results for the case $M/N = 3$, $N = 30$, $\xi = 3$ (with density $\rho = 0.87$). Panel (a): magnitude of the entries of the matrix $\mathbf{X}$. Panel (b): distribution of the modal energies of the original (square) and rotated (circles) basis functions. Panel (c): cosine of the angle between the modes of the original and rotated basis.

This appears to be a general trend since a similar behaviour was observed using the LASSO-based *a-posteriori* sparsification. A second major remark is that when the ratio M/N is increased, lower densities can be achieved for the same reconstructed energy whilst remaining within the feasibility region. Alternatively, rotating the original POD basis into subspaces of much larger dimension M gives increased flexibility to generate models capturing more energy for the same density.

6.1.4 Analysis of the rotated modal basis

We now move to the analysis of the rotated spatial and temporal basis functions (4.26). A system with $M/N = 3$, $N = 30$, obtained for $\xi = 3$, soon before the solution falls into the infeasible region in figure 6.3-(a), is analysed. Figure 6.5-(a) shows the magnitude of entries of the rotation matrix $\mathbf{X}$, found from the solution of (4.35). Figure 6.5-(b) shows the amount of kinetic energy $\tilde{\lambda}_i$ reconstructed by each mode on the new modal basis compared against the POD taken as reference. Figure 6.5-(c) shows the cosine of the angle between the old and new sets of spatial modes, defined as

$$\cos(\theta_i) = \frac{(\boldsymbol{\phi}_i, \tilde{\boldsymbol{\phi}}_i)}{||\boldsymbol{\phi}_i||_2||\tilde{\boldsymbol{\phi}}_i||_2}. \tag{6.2}$$

This quantity is also the diagonal of $\mathbf{X}$, since the relation $\mathbf{X}^\top\mathbf{X} = \mathbf{I}$ holds. A general

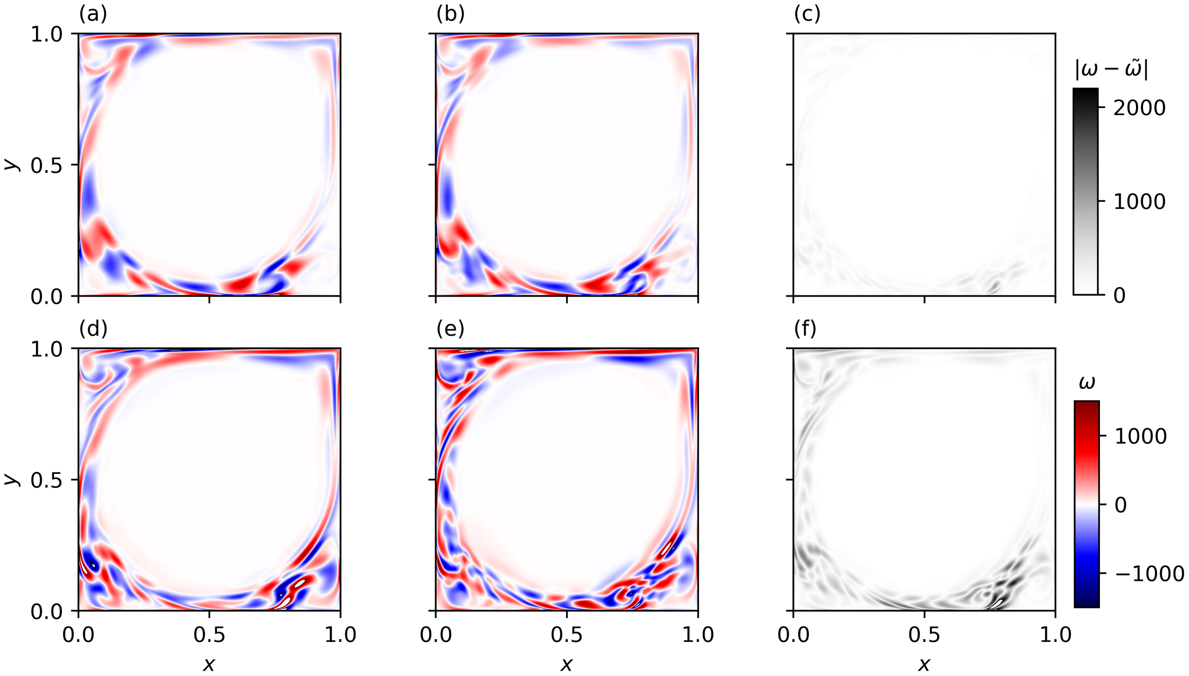

FIGURE 6.6: Vorticiy field for the firs POD mode, the first mode of the new basis and the absolute value of their difference in panel (a),(b) and (c), respectively. Panels (d),(e) and (f) same quantities for $i = 19$

trend can be observed. The high-energy modes are not significantly affected by the rotation while the opposite is true for the high-index, low-energy modes. More specifically, the first pair of eigenvalues, corresponding to the dominant fluid oscillation in the cavity, is virtually unchanged. This is then also observed for the first two spatial modes which are found to be almost collinear to the original ones, since $\cos(\theta_{(1,2)}) \sim 0.9$. A closer look at the rotation matrix $\mathbf{X}$ shows that, for the first two columns, the elements on the diagonal are close to 1 while all other entries are close to zero. Conversely, a different trend emerges when low-energy modes are considered. Figure 6.5-(c) shows that the angle between the original and rotated modes increases with the modal index. This behaviour derives from the problem formulation (4.35), constructed to generate a basis that minimises the energy loss with respect to the energetically optimal POD. As a result, the optimisation leaves unchanged the large and most energetic modes and rotates by a larger extent the less energetic modes to gain in sparsity. Figure 6.6 shows the out of plane vorticity component ω for the original POD mode $\boldsymbol{\phi}_i$, panels (a,d), and the rotated mode $\tilde{\boldsymbol{\phi}}_i$, panels (b,e), for indices $i = 1$ and 19, respectively. The third column, panels (c,f), shows the absolute value of the difference between original and rotated modes. As expected, the first mode is not considerably affected by $\mathbf{X}$. On the other hand, mode $i = 19$ is more pronouncedly affected by the rotation, with small scale features appearing along the shear layer. We argue that the introduction of small scale features is a combined effect of the stability constraint 4.35c, which enhances dissipation in the system (Balajewicz et al., 2016), and of the sparsity-promoting constraint (4.35b) since oscillatory modes are likely more effective in shrinking to zero the spatial

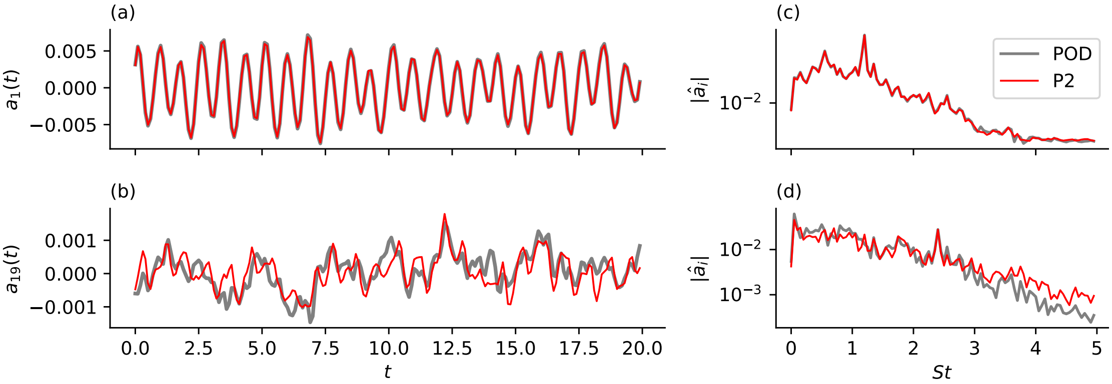

FIGURE 6.7: Temporal evolution and magnitude of the Fourier transform for $a_1(t)$ in panels (a) and (c), respectively, and for $a_{19}(t)$ in panels (b) and (d). Data is reported for the original POD temporal modes and for the rotated modes of the sparse system.

averages involved in the projection coefficients (4.9). The temporal modes as defined in (4.26) are affected similarly. This is illustrated in figure 6.7 showing the temporal evolution of modes $a_1(t)$ and $a_{19}(t)$, panels (a) and (b), and their amplitude spectra in panels (c) and (d), respectively. Since the first column of $\mathbf{X}$ is close to zero, except for X_{11}, mode $a_1(t)$ and its spectral content is not appreciably affected by the rotation, except for a small decrease of the amplitude due to decrease in energy content (see figure 6.5-(b)). Conversely, the spectral content of mode $a_{19}(t)$ is modulated by the rotation by introducing higher energy at high frequency components, consistent with the introduction of small scale features into the spatial mode.

6.1.5 Subset of interactions identified in the sparse model

Even if the modification to the temporal and spatial modes is small in relative terms, it is sufficient to introduce sparsity in coefficient tensor $\mathbf{Q}$ and the average triadic interaction tensor $\mathbf{N}$ when the domain integrals (4.9) and the temporal averages (4.16) are computed. To visualise how the sparsity varies when the penalisation weight ξ is increased, we introduce the tensor $\boldsymbol{\xi}$ with entries ξ_{ijk} defined as the value of the penalisation parameter ξ in (4.35) when the corresponding coefficient $\tilde{Q}_{ijk}$ is first shrunk to zero. Figure 6.8 shows two slices of $\boldsymbol{\xi}$ for $i = 1$ and $i = 30$, for the model obtained at $M/N = 3$ and $N = 30$. Panels (a) and (b) are obtained with the complete formulation P2 (including both the stability and sparsity promoting constraint), while panels (d) and (e), are obtained with the formulation P1 that only includes the sparsity promoting constraint. Panels (c) and (f) display the modal density ρ_i, defined as the density of the i-th ODE, as a function of the global density ρ, with ξ the curve parameter, for four modes across the hierarchy. For the complete formulation (P2), we observe that small-scale/small-scale interactions are the first interactions to disappear for moderate penalisation ξ. Increasing the penalisation, coefficients corresponding to interactions

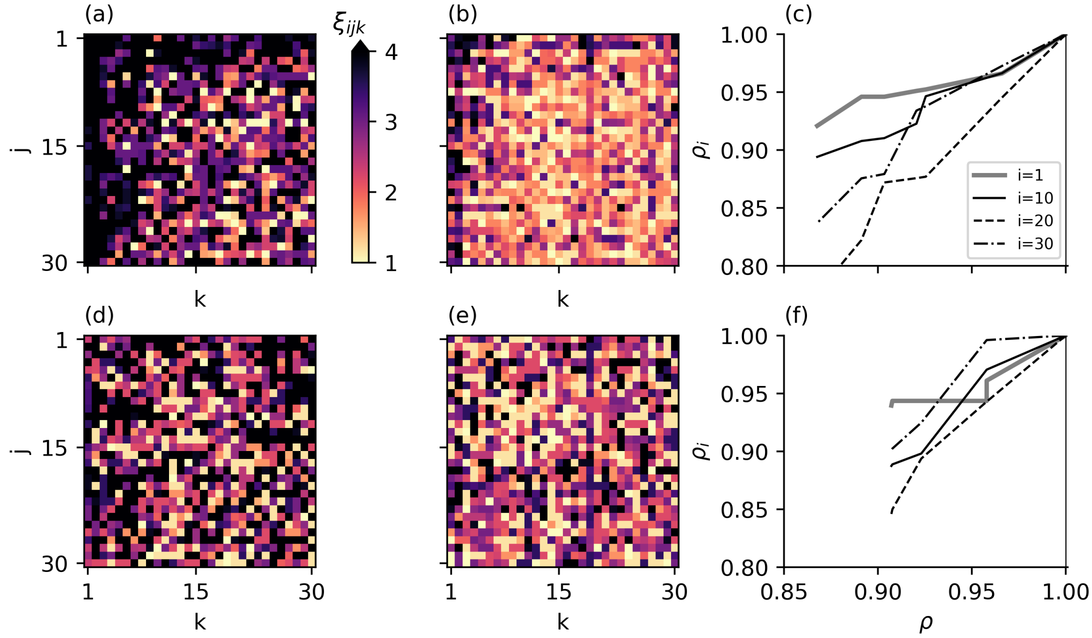

FIGURE 6.8: Distribution of the tensor ξ_{ijk} for $i = 1,30$ obtained solving the full problem P2 in panels (a) and (b), respectively. Panel (c): distribution of the curves $\rho_i - \rho$ for the solution of the full model. Distribution of the tensor ξ_{ijk} for $i = 1,30$ obtained solving the incomplete problem P1 in panels (d) and (e), respectively. Panel (f): distribution of the curves $\rho_i - \rho$ for the solution of the full model.

that are local in modal space are progressively pruned, leaving only coefficients capturing non-local interactions with the low-index modes for large penalisations. The structure of the sparse coefficient tensor obtained from the procedure is in agreement with the pattern of triadic energy interactions shown in figure 5.12. Interestingly, figure 6.8(c) indicates that high-index modes can be sparsified more efficiently. This might be related to the fact that high-index modes are also more modified during the optimisation while the low index modes are left almost unchanged to preserve the energy optimality imposed by (4.35a).

The key observation on the solution of the optimisation problem P1, bottom panels, is that the resulting sparsity pattern lacks the structure of the original triadic interaction tensor obtained with the complete formulation. In fact, even for small penalisations, coefficients corresponding to energy-relevant interactions are shrunk to zero. A random, chessboard-like coefficient distribution appears, even though similar global density is observed with similar penalisation ξ. It is arguable that this is due to the fact that introducing a penalisation on the rotated tensor $\tilde{\mathbf{Q}}$ without any constraint on the temporal evolution of the rotated system leads to a uniform reduction of the coefficients $\tilde{Q}_{ijk}$ across all modes, because the strength of the energy interactions are not taken into account. Instead, since the coefficients Q_{ijk} are initially nonzero and do not have a clear structure (see figure 5.12), a uniform decrease of their magnitude across all modes will result in a chessboard-like pattern of the retained interactions. This behaviour can also be noticed in the shape of the $\rho_i - \rho$ curves, panel (f), where modes are almost equally

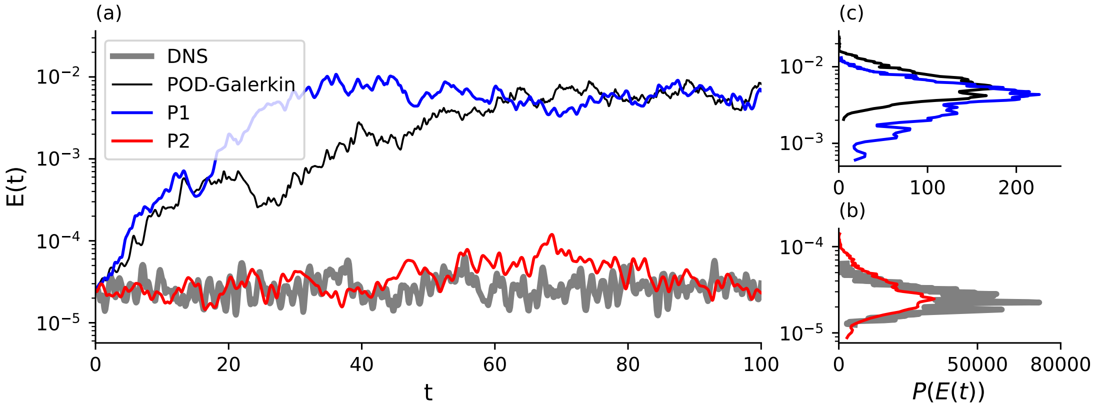

FIGURE 6.9: Temporal evolution of the turbulent kinetic energy $E(t)$. The performance of the two sparse models obtained by solving problem P1 and P2 are compared against the DNS and the dense model obtained with Galerkin projection. Panels-(b,c) show the probability distribution of $E(t)$.

sparsified. This behaviour is a consequence that, for the two formulations P1 and P2 of the problem, two different minima with similar values of the minimal density but with a different shape of the retained interactions have been found.

6.1.6 Temporal integration and energy analysis of the sparsified system

In this section, we analyse the properties of the sparse reduced-order model under temporal integration and consider the triadic interactions tensor **N** and the average energy budget (see equation (4.15)). The same configuration studied in previous sections, with $M/N = 3$, $N = 30$ and $\xi = 3$, is considered. Reduced-order models obtained by solving problems P1 and P2 are considered. These models are integrated forward in time from an initial condition with an implicit time stepping scheme for a time-span equal to $T = 500$ time units, from an initial condition obtained from one of the snapshots utilised for the POD. Figure 6.9-(a) shows the first one hundred time unity of the temporal evolution of the turbulent kinetic energy (4.15) for these two models, compared with the evolution from DNS and from the original POD-Galerkin model. Figure 6.9-(b,c) show the probability density function of the same quantity. As expected, the kinetic energy from the POD-Galerkin model simulation rapidly overshoots and converges towards energy levels about two orders of magnitude larger than the reference value from DNS. This behaviour is a well known result (see e.g. Noack et al. (2016); Östh et al. (2014)). Similarly, the sparse model obtained without any constraint on the temporal stability (problem P1) reproduce this behaviour. This can also be seen in the probability distribution functions in panels (b) and (c) of the turbulent kinetic energy showing that the two former models predict a larger value of the turbulent fluctuations. Conversely, the prediction of the sparsified and stabilised model (problem P2) matches well the prediction of the DNS both as average value and amplitude of the fluctuations.

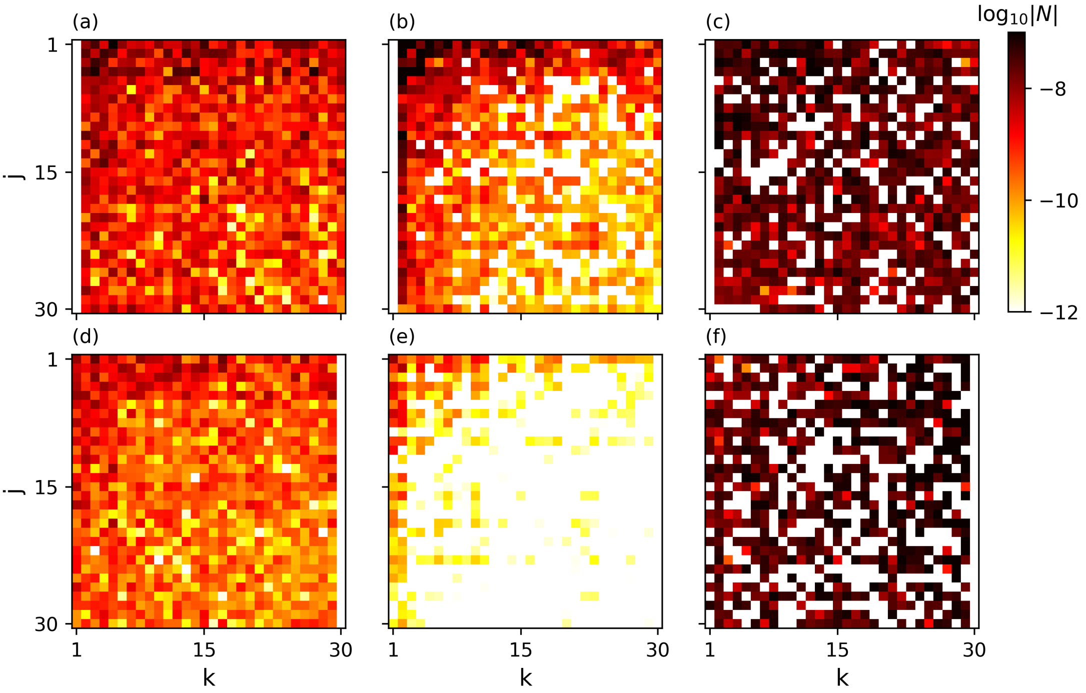

FIGURE 6.10: Intensity of the average nonlinear energy transfer rate $\tilde{\mathbf{N}}$ for mode $i = 1$ in models obtained from projection and formulations P2 and P1 in panels (a), (b) and (c), respectively. Panels (d), (e) and (f) show the same quantity for mode $i = 30$.

This result indicates the need for a tie back to the temporal evolution of the system to obtain meaningful results.

Once obtained the new sets of temporal modes, the triadic interactions tensor $\mathbf{N}$ and linear term $\sum_{i=1}^{N} \tilde{L}_{ij}\tilde{a}_i\tilde{a}_j$ in (4.15) can be computed. The structure of the nonlinear energy transfer rate tensor $\tilde{\mathbf{N}}$ for the rotated Galerkin models is reported in figure 6.10, showing two slices for $i = 1, 30$ (first and second row, respectively). Panels (a) and (d) refer to the model obtained from Galerkin projection onto the original POD subspace and will be used as reference. Here the temporal coefficients are the projections onto the DNS snapshots. Data for the model obtained from problem P2 is shown in panels (b) and (e) while data for the model obtained with the P1 formulation is shown in panels (c) and (f). For these models, temporal coefficients are obtained from the forward integration of the Galerkin models. It can be observed that the model from formulation P2 displays a pattern of interactions consistent with the DNS data projected onto the original POD modes, both in terms of the organisation across modes as well as in strength. The organisation of the interactions is similar to that in the dense Galerkin model obtained using the *a-priori* stabilisation method of (Balajewicz et al., 2013), although a more aggressive reduction in the strength of the interactions between the high-index modes is observed. By contrast, energy interactions in the model obtained from the formulation P1 are orders of magnitude more intense, across all triads. This is a manifestation

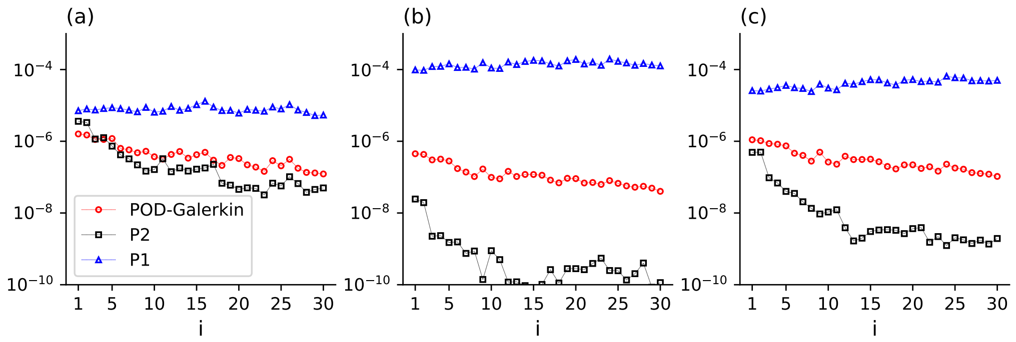

FIGURE 6.11: Value of the sum of the absolute value of N_{ijk} for the different models in the three regions of **N** defined in figure 5.12-(a). Panel (a),(b) and (c) show the regions LL,HH and LH, respectively.

of the lack of dissipation and the consequential over-prediction of energy, across all modes. Interestingly, it can be observed that the HL/LH asymmetry observed in figure 5.13 is preserved throughout the sparsification. Conversely, this physical feature is lost in the *a posteriori* approach in the construction of the database matrix to preserve the mathematical consistency on the regression problem. Arguably, the *a priori* approach is preferable for systems with a complex structure of **N**.

To visualise the overall structure of the tensor **N** across the hierarchy of modes $i = 1, \ldots, N$, we apply the methodology explained in the energy analysis section. To characterise which regions of **N** are more active in the average energy transfer, the sum of the absolute values of the entries N_{ijk} for the four different areas LL,HH and LS as defined in figure 5.12-(a) is computed. Figure 6.11 shows the value of LL,HH and LH in panel (a),(b) and (c), respectively. Different trends for the two models P1 and P2 with respect to the POD-Galerkin model are observed. First, we consider panel (a) containing the LL interactions. We observe that the value provided by the P2 formulation is very similar to the projection model obtained by DNS with the only difference of a slightly lower value due to the presence of the constraint on **Q**. In addition, the model obtained with the P2 approach has the same decay as the projection model (DNS). Conversely, the P1 model does not present any decay of the energy level across the hierarchy of modes. This trend of the model obtained with the P1 approach is shown also in the interactions of the region HH and LH, in panels (b) and (c), respectively. This is consequence of the over prediction in the amplitude of the temporal fluctuations especially observed for the modes with high modal index i. In addition, in panels (b) and (c) it can be observed that, on average, the value of the model P2 is lower than the one obtained by projection, this effect is much weaker with respect to the interaction of the region LL. This confirms a consistent sparsification where the less relevant interactions, mainly present in the regions HH and HL/LH, are more heavily pruned than interactions present in the region LL. More interestingly, we observe that the numerical difference between $P2$ and the DNS in panel (b) and (c) grows as we move towards

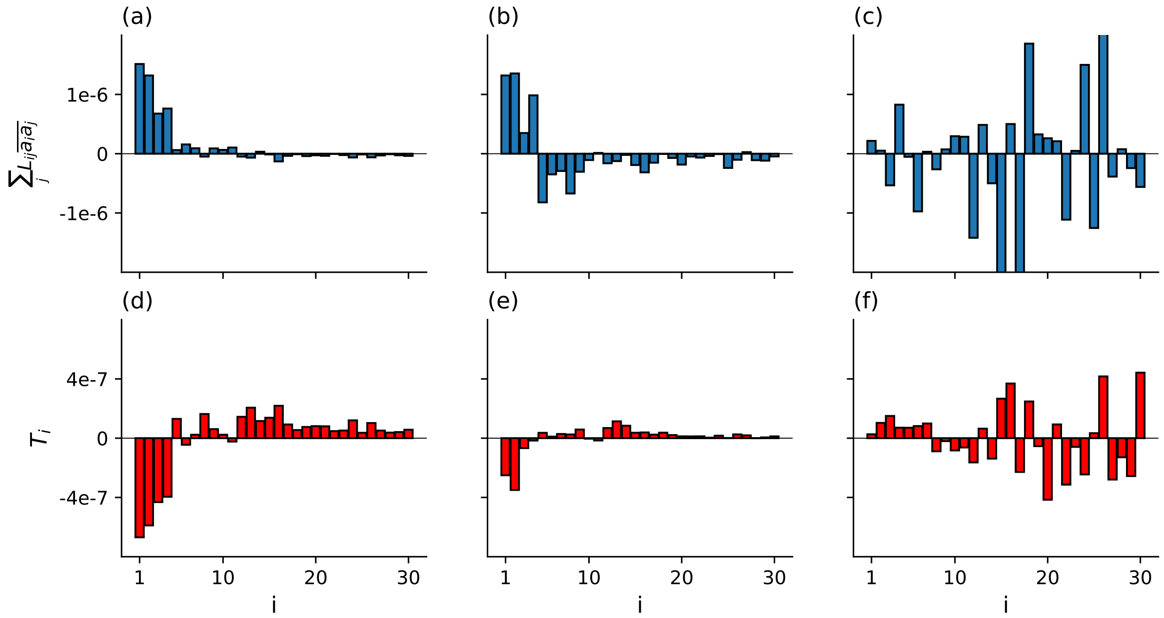

FIGURE 6.12: Value of the linear term in (4.15) for the Galerkin projection, the P2 and the P1 approach in panels (a),(b) and (c), respectively. Value of the quadratic term T_i in (4.15) for the Galerkin projection, the P2 and the P1 approach in panels (d),(e) and (f), respectively.

larger values of i showing that high-index modes are more effectively sparsified as already shown in figure 6.8-(c). To conclude we observe that none of these effects are present for the P1 model confirming that this approach is not suitable for a physically consistent sparsification.

We now analyse mode by mode the linear and quadratic energy transfer terms as defined in the mean average energy balance (4.15). In figure 6.12 we consider three different systems: the one obtained by Galerkin projection onto the POD subspace in panels (a) and (d), the one obtained solving the problem P2 in panels (b) and (e) and solving problem P1 in panels (c) and (f). The model obtained by projection is analysed first. In panel (a) the linear term in (4.15) obtained by POD-Galerkin projection is shown. We observe a large production of kinetic energy in the first modes and a moderate dissipation in the last ones. Overall this results in an unbalance between production and dissipation leading to the over prediction of kinetic energy shown in figure 6.9-(a). The non linear term T_i shown in panel (d), instead, displays the correct energy cascade from the large (negative values of for small i) to the small scales (positive values for i large) already observed for two dimensional turbulent flows. Second, we consider a model obtained by the solution of the complete optimisation problem (P2) in panels (b) and (e). Panel (b) shows that for this case the production/dissipation is better balanced with lower production in the low index modal structure and larger dissipation in the high index modes. In addition, the P2 approach also manages to preserve the right shape of the nonlinear energy transfer between modes as displayed in panel (e).

Conversely, the absence of the constraint in the optimisation problem (formulation P1) does not provide any balance of the turbulent kinetic energy budget (4.15). Both the linear and the quadratic terms are off with respect to the value and the shape obtained from the POD-Galerkin model shown in panels (a) and (d). In addition, the nonlinear energy transfer appears to be uniformly diminished and reshuffled losing the physics of the corresponding modal decomposition as shown in panel (f). The largest amplitude of T_i for the high index modes is a consequence of the unphysically large value of the fluctuations in the high index temporal coefficients. These results show that to obtain a sparse model that is consistent with the physics of the flow it is of paramount importance to introduce in the sparsification algorithm some information about the dynamics or the physics of the flow. This is consistent with what was observed by Loiseau and Brunton (2018) that proved the need for additional constraints on the energy conservation to en enhance the temporal accuracy of dynamical system reconstructed from data. These findings are in agreement with the results obtained for the *a posteriori* sparsification framework confirming that including the temporal dynamics $\dot{a}_i$ in the sparsification preserves the correct energy transfer in the sparsified system.

6.1.7 Temporal stability *a posteriori* vs *a priori*

In this section, the long term temporal stability of the two different sparsification approaches is discussed. We consider two models of the lid driven cavity flow, both resolving 95% of the kinetic energy, obtained with the *a posteriori* and the *a priori* approach, respectively. For the *a posteriori* approach the optimally sparse system has been identified with the l_1 pseudo-density criterion explained previously. On the other hand, in the case of the *a priori* sparsified approach the optimally sparse system has been chosen as the sparser system still satisfying all the optimisation constraints. These choices are not unique but it is in the author opinion that different choices do not change sensibly the results. Figure 6.13 was generated by integrating the optimally sparse system from random initial conditions for a time interval three times longer than the DNS simulation used for generating the decomposition and training the sparsification algorithms. Panels (a) and (b) display an *a posteriori* and an *a priori* model, respectively. The red line is the DNS simulation, used as a reference, while the other continuous lines represent different initial conditions randomly generated. Panel (a) shows results of the time integration of the *a posteriori* sparsified models. Interestingly, two trends emerge. First, the turbulent kinetic energy seems to keep memory of the initial condition. This is shown in panel (a) where the blue line, starting from a slightly different initial condition, although having the correct fluctuation amplitude, keeps the memory of the initial condition predicting the wrong level of turbulent kinetic energy. Second,

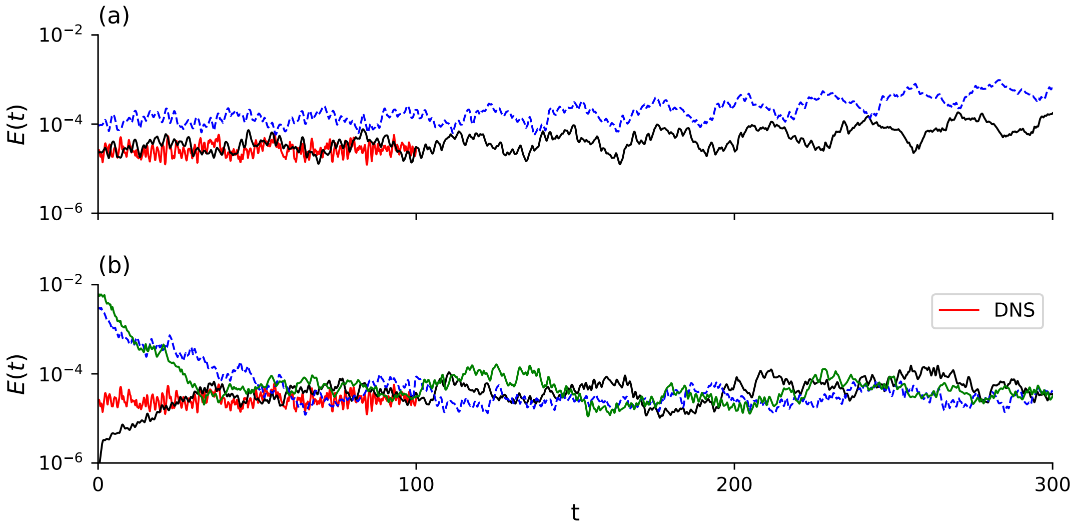

FIGURE 6.13: Effect of a variation of the initial condition of the long time temporal integration of the sparse system. Panel (a) and (b) show results for the *a posteriori* and *a priori* sparsified systems, respectively. The reference values of the DNS is displayed as a red line while simulations obtained with different values of the initial conditions are shown as continuous black,blue and green lines, respectively.

both sparse solutions (identified by the black and blue lines) tend to drift away from the correct energy level towards higher ones showing the inception of a long term temporal instability for time integration intervals longer than the training time-span ($T = 100$). This was expected since the model is trained on the data coming from DNS (of length $T = 100$), therefore no stability outside the training dataset can be guaranteed. In addition, it can be observed that as the flow evolves some small scales oscillations are smeared away, this behaviour is particularly evident for $T > 250$ times unit where only the large time scales oscillations are retained. It is arguable that this is a side effect of the sparsification performed with constant regularisation weight γ_i that tends to penalise more the low amplitude fluctuations. Interestingly, both of these behaviours are not present for the *a priori* sparsified systems shown in panel (b). It can be observed that all the simulations, even if starting from different initial conditions after an initial transient converges towards the correct level of turbulent fluctuations without any long term stability issue. In addition, no loss of small scales fluctuations is observed in this model even for times outside the training region. This result shows that the *a priori* sparsification technique provide superior performances both for the short and long time integration of the sparse model. The price to pay is a higher computational cost of the sparsification procedure and generally the impossibility to reach densities as low as for the *a posteriori*. This behaviour is consequence of the mathematical formulation of the two approaches. In fact, for the *a posteriori* approach the temporal stability can be ensured only for times interval equal to the length of $\dot{\mathbf{A}}_i$. Since this quantity is computed projecting the Eulerian acceleration field onto the spatial modes it is available only for a time span comparable with the length of the numerical simulation. When

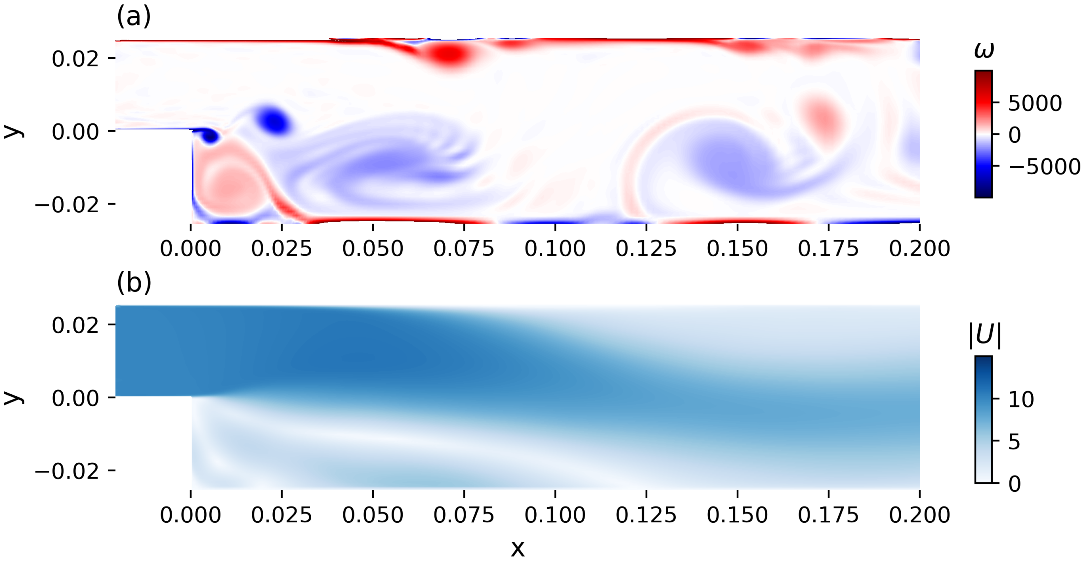

FIGURE 6.14: Instantaneous vorticity field and average of the magnitude of the velocity field in panels (a) and (b), respectively.

complex flows, requiring a large number of snapshots to reach statistical convergence, are considered, the computation of this quantity can easily become challenging. Conversely, for the *a priori* approach, in the constraint (4.35b) there is no limit on the length of the time interval over which the mean of the kinetic energy can be computed. Since the integration of the sparse system has a low computational cost it is possible to extend the stability for much longer time intervals with respect the length of the data needed to obtain the POD.

6.2 2D backward facing step

In this section the *a priori* sparsification technique is applied to the two dimensional flow evolving over a backward-facing step. The domain is defined in Cartesian coordinates $\mathbf{x} = (x, y)$ and the velocity field is defined by the components $\mathbf{u} = (u, v)$. The Reynolds number is defined as $Re = HU/\nu$, where H, U and ν are the step height, the inlet velocity and the kinematic viscosity, respectively, is set equal to 2×10^4. The DNS has been performed in OpenFOAM with the icoFOAM solver, without any turbulence model. Conversely to the cavity flow, there is no need to save the Eulerian acceleration filed at each time step since the *a priori* framework does not need this quantity as input. Therefore, no modification of the solver was deemed necessary. Figure 6.14 shows the instantaneous out of plane vorticity field ω for a snapshot at a given instant of time and the magnitude of the average velocity field, defined in terms of the Cartesian coordinates as $U = \sqrt{u^2 + v^2}$, in panel (a) and (b), respectively. In panel (b) we observe a large re-circulation area downstream the step and near the upper wall characterised by the magnitude of the velocity field reaching zero value. Panel (a) shows these flow features being strongly linked to the generation of vortices in the shear layer near the corner of the step (blue flow structures) and on the top wall in correspondence to the deceleration of the flow (red flow structures in panel (a)). This flow has higher complexity with respect to the cavity flow previously analysed. Thus, more modes will be required to construct reduced order models describing a meaningful amount of spatial scales. In addition, the open inlet and outlet boundary conditions make the pressure term in the Galerkin models not identically zero. Therefore, the sparsification framework must balance also the effects of neglecting this term making the generation of accurate reduced-order models challenging. Here we considered an initial POD-based reduced-order model composed by $N = 50$ POD modes, resolving up to 95% of the turbulent kinetic energy. Similarly to what was done for the NACA test case the pressure term, different from zero, is not included in the generation of the ROM (4.6). The idea is to leave to the sparsification framework, in particular to the stability constraint (4.35c) the task of tuning the model coefficients such that the effects of the pressure term are taken into consideration. Leveraging both our observations on the sparsification of the two dimensional lid-driven cavity flow and the results of Balajewicz et al. (2013) about the stabilisation of reduced order systems, a ratio $M/N = 2$ has been chosen, corresponding to a dimension $M = 100$ for the larger subspace. The two values $N = 50$, $M = 100$ corresponds to the limit of the computational resources used in this study.

6.2.1 Energy analysis

We now analyse the properties of the tensors $\mathbf{N}$ and $\mathbf{Q}$, using as a reference, the dense model obtained by Galerkin projection of $N = 50$ POD modes. Figure 6.15 shows the base-ten logarithm of the absolute value of the entries N_{ijk} for $i = 1{,}10{,}50$. In addition, panel (a) shows the subdivision of the interactions governing the dynamics of the i-th mode into the four regions already defined in figure 5.12-(a). Namely, LL contains the interactions of large modes between themselves, LH, HL the mixed interactions between large and small structures and HH the interactions of high index modes. All the interactions are generally non-zero, this results from the conjoint effect of the tensor $\mathbf{Q}$ being generally non-zero and the temporal coefficients $a_i(t)$ having a complex spectral structure. However, a subset of dominant interaction exists. Namely, the interactions belonging to the regions LL,HL and LH are generally stronger with respect to the ones of the HH region. Similarly to what was discussed in the lid-driven cavity, this results from the strong advective nature of the flow in consideration, where small structures are transported downstream by larger and more energetic structures. For better quantify this effect on the triadic interactions, figure 6.16 describes quantitatively the asymmetry of the average triadic interactions tensor. Panel (a) displays the value of the function χ, as previously defined in (5.12), as a function of the normalised bandwidth n/N as defined in figure 6.15-(a). Panel (a) show that $\chi > 1$ for every n/N

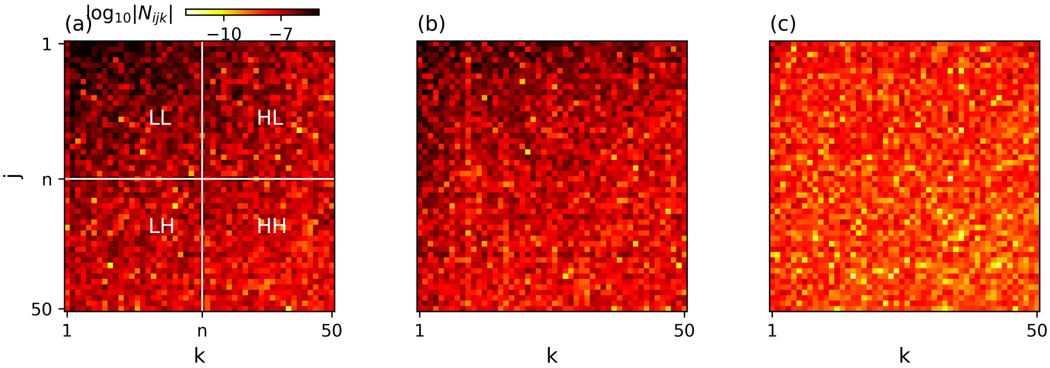

FIGURE 6.15: Magnitude of the average triadic interactions tensor **N** for $i = 1, 10$ and 50 in panel (a),(b) and (c), respectively. In addition, panel (a) shows the four different regions of **N** corresponding to interactions of different nature.

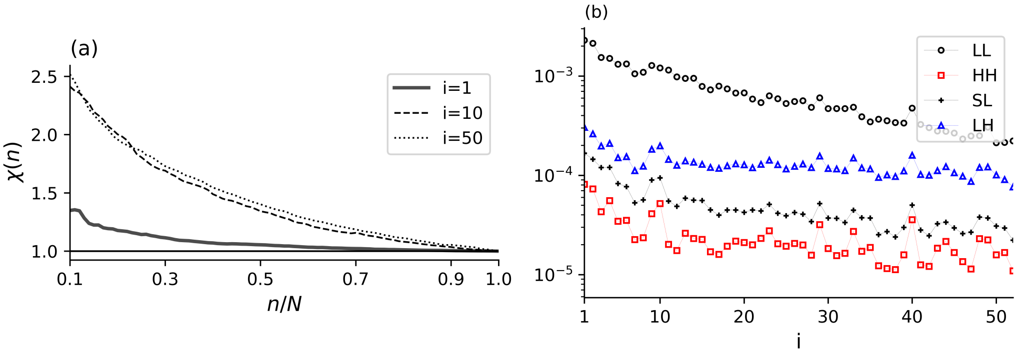

FIGURE 6.16: Panel (a) shows the coefficient $\chi(n)$ as a function of the normalised cutoff n/N for the same three modes considered in in figure 6.15. Panel (b) shows the sum of the entries of **N** for the four regions defined in figure 6.15 as a function of the modal index i.

and converge to 1 as $n/N \to 1$. This is consequence of the fact that the region HL is predominant with respect to the region LH. From a physical standpoint, this means that the flow is characterised by a strong advection of the small structures by the large ones. In addition, we observe that χ grows as the modal indexes i grows. This result is expected and is consequence that as the modal index i grows the corresponding mode describes spatially smaller structures where the advective effect is stronger. To characterise energy paths across all the structures we compute the sum of the absolute value of the entries N_{ijk} contained in the four regions defined in 6.15-(a) for each slice of N_{ijk} corresponding to a modal index. The result is shown in figure 6.16-(b). First, it can be observed that energy transfers in the regions LL, HL and LH are generally more intense than those in the region HH. This follows from the observation that large-scale/large-scale and large-scale/small-scale interactions are more relevant with respect to the small-scale/small-scale interactions, across the entire hierarchy. In addition, we observe that the interactions LH are always larger than the HL interactions. This is consequence of the asymmetry in the tensor **Q** and not of the temporal coefficients. This asymmetry is in agreement with the picture of energy transfers between

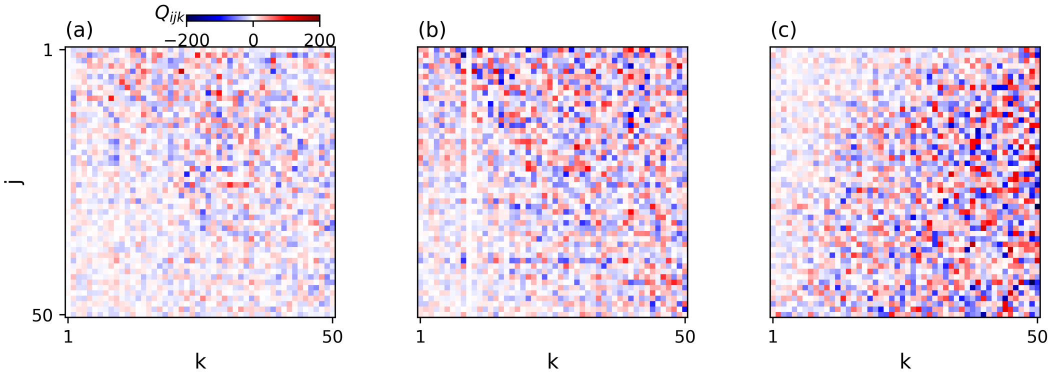

FIGURE 6.17: Maps of the entries of the tensor $\mathbf{Q}$, for $i = 1$,10 and 50,

scales in homogeneous isotropic two-dimensional turbulence (Ohkitani, 1990; Laval et al., 1999), where the large scales interact with the small ones in a non-local fashion. Lastly, it can be observed that the structure of the energetic interactions is mainly determined by the temporal coefficients as the tensor $\mathbf{Q}$ constructed from projection modes does not possess any sparsity structure. Figure 6.17 shows the entries of $\mathbf{Q}$ for $i = 1$,10 and 50 in panel (a),(b) and (c), respectively. We observe that all the coefficients Q_{ijk} are generally different from zero and no sparsity structure is present. Interestingly, despite being all different from zero the entries Q_{ijk} have a structure driven by the advective nature of the flow. However, it is important to stress that due to the wider range of numerical values spanned by $\overline{a_i a_j a_k}$ with respect to Q_{ijk} the shape and the relative magnitude of the average triadic interactions in driven mainly by the temporal coefficients.

6.2.2 *A priori* Sparsification

Here we applied the *a priori* sparsification procedure to a model with $N = 50$ generated by the rotation of the POD modes in a subspace of larger dimension $M = 100$. Because of the definition of the optimisation problem (4.35) a family of different models can be obtained increasing progressively ξ starting from 1 (where no sparsification is promoted) to the maximum value of ξ where the solution exits the feasibility region of the problem. This value of ξ corresponds to the sparser system for two give values of M and N. In addition, for every value of ξ an external loop to find η to ensure long time accuracy of the model has been added. For every value of ξ the density ρ and the energy resolution e_N is computed and displayed in figure 6.18. Figure 6.18-(a) shows the $\rho - e_N$ plane for the model considered. As previously observed for the cavity flow the rightmost point $\rho = 1$ loses some ability of reconstructing energy with respect the original basis. This is consequence of the stabilisation constraint that slightly modifies the POD drifting away the new basis from the energetic optimal POD. Further increasing ξ, we observe that lower values of ρ can be reached. Increasing ξ further, $\rho \sim 0.88$ is reached, here the solution of (4.35) falls outside the feasibility region, i.e. it is not possible to find a rotation $\mathbf{X}$ that decreases further the density and, at the same time, satisfies

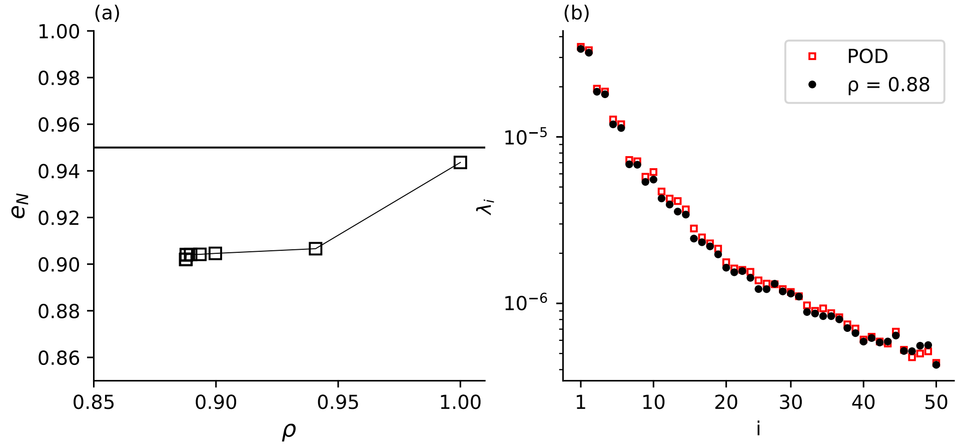

FIGURE 6.18: Panel (a): $\rho - e_N$ curve for a model constructed by $N = 50$ modes. Panel (b): POD eigenvalues λ_i for the system with the minimal density $\rho = 0.88$ (leftmost point in panel (a)) compared with the original POD basis.

both the stability and sparsity constraints. Observing the shape of the $\rho - e_N$ curves, it can be observed that after a sudden decrease of e_N it is possible to promote sparsification without affecting considerably the ability of the new basis of reconstructing kinetic energy. It is arguable that this is consequence to higher physical complexity of the flow considered with respect the lid driven cavity flow. Similarly to what done for the lid-driven cavity flow, since all the systems obtained have good long term stability properties we select as sparse system the leftmost system in the $\rho - e_N$ plane shown in panel (a). Panel-(b) shows the eigenvalues λ_i of the sparse decomposition against the POD ones. Similarly to what observed for the cavity flow the first pairs of eigenvalues are only marginally affected, since the optimisation problem aims to the energetic optimality with respect to the original POD. Conversely as we move to higher modal indexes the eigenvalues are more heavily modified.

6.2.3 Interactions identified by the sparsification

The modification of of the temporal and spatial modes induced by $\mathbf{X}$ even if small is sufficient to introduce sparsity in the rotated quadratic interaction tensor $\mathbf{Q}$ and consequently in $\mathbf{N}$. To better visualise the sparsity pattern as the regularisation weight ξ increases, we introduce the tensor $\boldsymbol{\xi}$ with entries ξ_{ijk} defined as the value of the penalisation parameter ξ in (4.35) when the corresponding coefficient $\tilde{Q}_{ijk}$ is first shrunk to zero. Figure 6.19 shows three slices of $\boldsymbol{\xi}$ for $i = 1,10$ and 50. We observe that small-scale/small-scale interactions (HH region) are the first interactions to disappear for

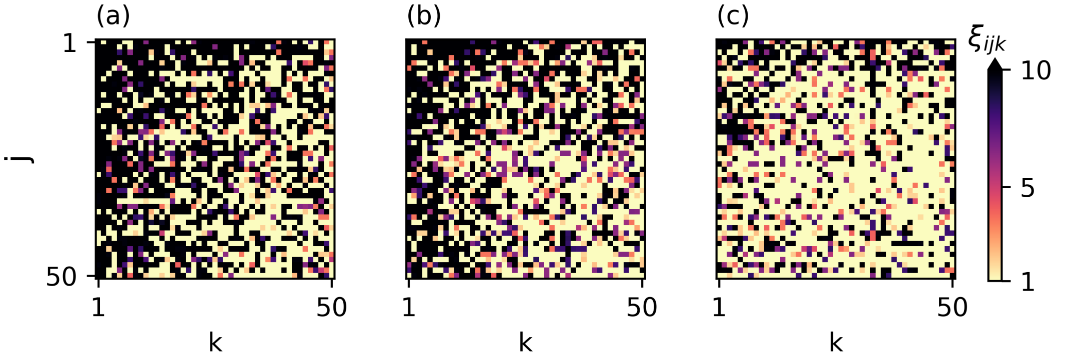

FIGURE 6.19: Distribution of ξ for $i = 1, 10$ and 50 in panel (a), (b) and (c), respectively.

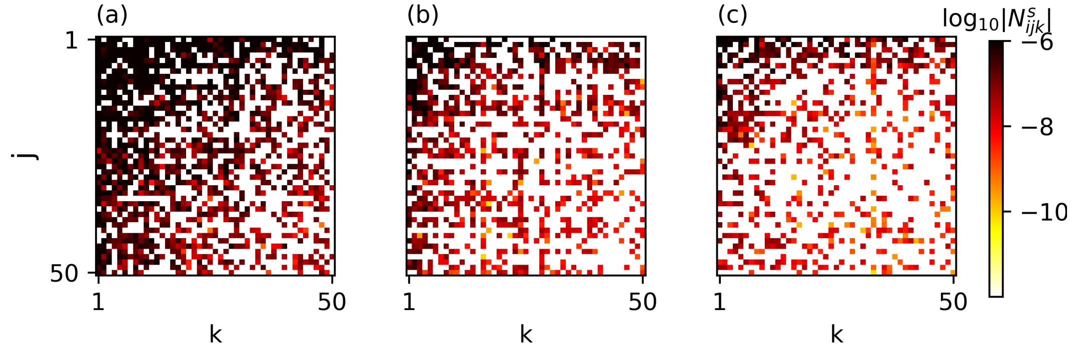

FIGURE 6.20: Distribution of the sparse interaction tensor $\mathbf{N}^s$ for $i = 1{,}10$ and 50 in panel (a),(b) and (c), respectively.

moderate penalisation ξ. Increasing the penalisation, coefficients corresponding to interactions that are local in modal space are progressively pruned, leaving only coefficients capturing non-local interactions with the low-index modes for large penalisations. The structure of the sparse coefficient tensor obtained from the procedure is in agreement with the pattern of triadic energy interactions shown in figure 6.15. More interestingly, the non-symmetry between the HL and LH region is preserved, this is particularly evident for the mode $i = 50$ figure 6.19-(c). In figure 6.20 the sparse triadic interaction tensor $\mathbf{N}^s$ computed with the rotated modes is shown for the same three modes considered in figure 6.19. It can be observed that the sparsified model has a pattern of interactions resembling that of the dense model in figure 6.15. However, weak interactions and the associated flow physics have been pruned. It is also clear that the asymmetry of the interaction pattern observed in figure 6.15 and the physical mechanism that originates is preserved by the *a priori* sparsification. This structure can be more completely visualised in figure 6.21. Panel-(a) displays the structure function χ, as defined in (5.12), against the non-dimensional bandwidth n/N. The correct trend observed for the POD reference case shown in figure 6.16-(a) is also preserved in the sparse system. Panel-(b) shows the sum of the absolute values of the four different

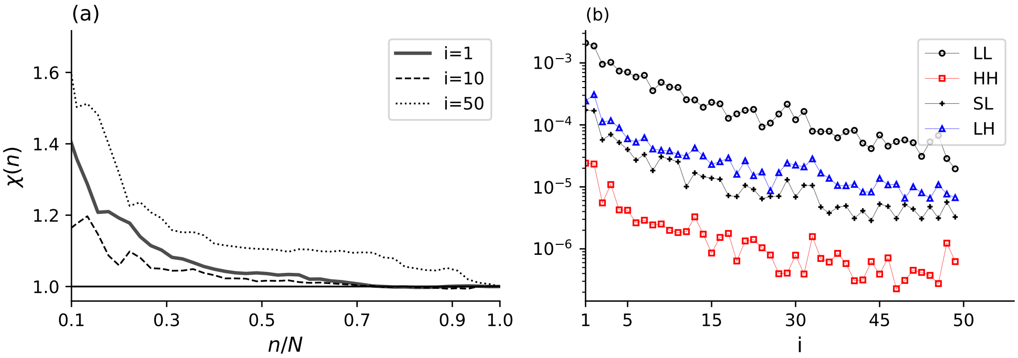

FIGURE 6.21: Panel (a): coefficient $\chi(n)$ as a function of the normalised cutoff n/N for the same three modes considered in in figure 6.18. Panel (b): sum of the entries of $\mathbf{N}$ for the four regions defined in figure 6.18 as a function of the modal index i. All the results are computed for the sparse triadic interaction tensor.

regions of N^s_{ijk} against the modal index i. Qualitatively, the behaviour observed in figure 6.16-(b) is preserved and the relation LH $>$ HL is true for every i. In addition, it can be observed that HH are smaller than the original ones as direct consequence to the fact that the sparsification affects more predominantly the interactions belonging in this area.

6.2.4 Time integration and flow field reconstruction

Once the optimisation problem (4.35) is solved a new set of matrix coefficients are obtained from (4.27). Here we turn our attention to temporal integration of the sparse Galerkin system defined in terms of this new modal basis. The model composed by $N = 50$ and $M/N = 2$ is integrated for twice the time-span available from DNS, *i.e* $T = 2$ non dimensional time units. Figure 6.22 displays the time evolution of the predicted turbulent kinetic energy $E(t)$ (4.13) in red, against the reference value obtained by DNS displayed as thick grey line in the background. The sparse model is able to reproduce both the average level and the fluctuation of the turbulent kinetic energy. More interestingly, the sparse system has good long term accuracy maintaining the correct level of turbulent fluctuations for longer time interval than the data-set used to generate the modes. No upwards drift is present as observed for the *a posteriori* sparsification of the cavity flow in figure 6.13. Arguably, the better long time performances are due to the conservation of the shape of triadic interactions by the *a priori* sparsification methodology. Since this particular shape of the triadic interactions can be related to advective dynamics of the flow, a good conservation of the shape of triadic interactions in modal space is crucial to obtain model faithfully describing the advective flows.

Once the new temporal $\tilde{a}_i(t)$ and spatial $\tilde{\boldsymbol{\phi}}_i$ coefficients are known from (4.26) the flow field can be reconstructed according to (4.3). Figure 6.23 shows the longitudinal and

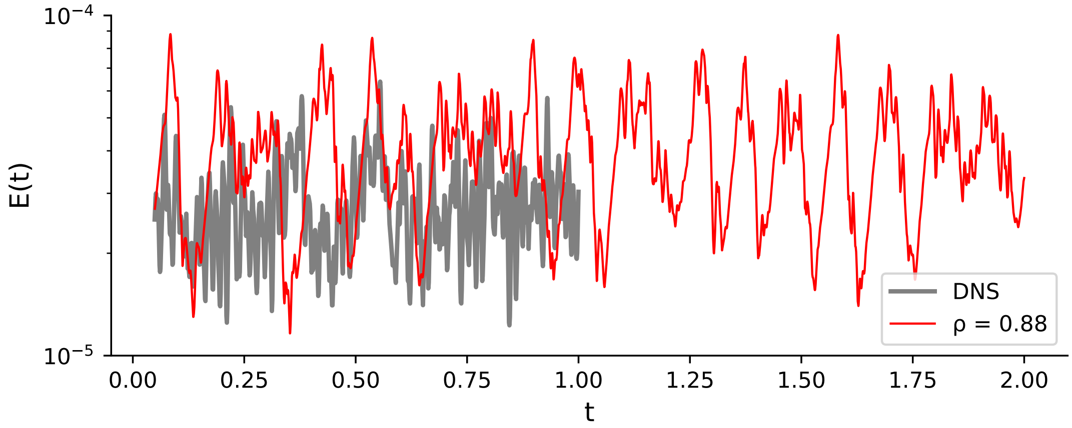

FIGURE 6.22: Temporal evolution of the turbulent kinetic energy $E(t)$, reference value from DNS against the value predicted by the sparse model.

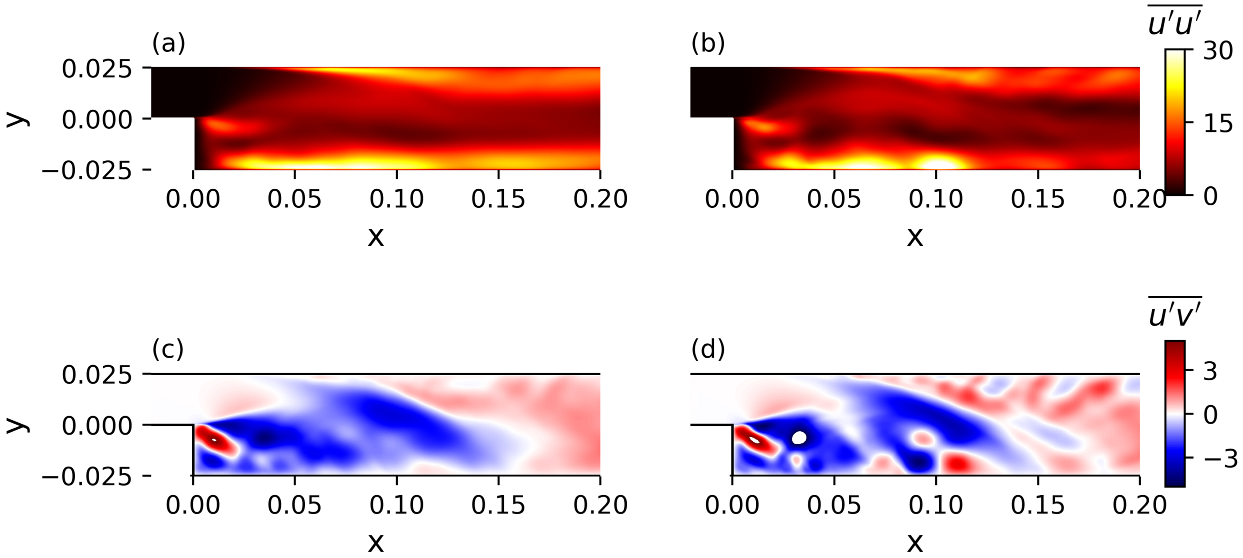

FIGURE 6.23: $\overline{u'u'}$ component of the Reynolds stresses obtained from DNS and from the sparse system in panel (a) ad (b), respectively. $\overline{u'v'}$ component of the Reynolds stresses obtained from DNS and from the sparse system in panel (c) ad (d), respectively.

traversal Reynolds stresses. The first column, panels (a) and (c), shows the field computed by DNS, while panels (b) and (d) display the results obtained from the temporal integration of the sparse system. Overall, we observe that the global structure of the Reynolds stresses is well predicted by the sparse system both in the spatial distribution and in the numerical value. More specifically, the sparse model is able to reproduce the correct Reynolds stresses downstream the step and in the boundary layer in the bottom and top wall near the separation and the reattachment of the boundary layer. The only difference appears to be in $\overline{u'v'}$ near the top wall where some fluctuations appear. We argue that this is consequence to the fact that since the optimisation problem (4.35) tries to preserve the original energetic optimality of the POD, the small scales are modified more aggressively. This leads to a less accurate description of the evolution of the small scales with respect to the large ones.

6.3 Summary: *a priori* sparsification

This final section aims to collect the major findings obtained from the *a priori* sparsification framework. To begin with, we considered the two-dimensional lid-driven cavity flow at Reynolds number $Re = 2 \times 10^4$. The cavity flow is a classical test case already utilised for the *a posteriori* framework allowing a comparison between the two methodologies to better understand their performances. In addition, we consider the evolution of the separated flow over a backward facing step. Similarly to what done for the *a posteriori* framework, sparse models are generated by varying the regularisation weight ξ in the sparsity promoting constraint (4.35b). Subsequently, different models are visualised onto the reconstructed energy - model density plane. The major outcomes of this analysis are the followings.

- The model size N and the ratio M/N affect the energy reconstruction - sparsity trade off. Larger models can be sparsified more with less impact on the energy reconstruction ability. In addition, lower densities can be obtained by rotating the original POD basis within larger sub-spaces (a higher ratio M/N).
- The formulation of the sparsity constraint allows the definition of a minimal reachable density for given values of the parameters N and M/N. This is a clear advantage with respect to regression based methods, where the regularisation weight is a free parameter that needs to be carefully chosen by the user after the evaluation of the performances of the single models.
- From a more physical perspective, it can be observed that the structure of the inter-modal energy interactions in the rotated model is qualitatively similar to that of the original POD-Galerkin system. More specifically, coefficients of the quadratic interaction tensor corresponding to large-scale/large-scale and the large-scale/small-scale interactions are preserved, but those defining the weak small-scale/small-scale interactions are pruned during the optimisation. This behaviour was observed for both the test cases analysed.
- Interestingly, it was also observed that a physically-consistent distribution of the interactions and a stable long-term behaviour can only be obtained by augmenting the sparsification procedure with a temporal stability constraint. In fact, it was observed that models obtained without such constraint can inherit well-known temporal stability issues of dense POD-Galerkin model.
- Once the stability constraint has been included in the formulation, this approach shows better prediction abilities both on short and long time spans with respect to the regression based *a posteriori* approach. In addition, good robustness with respect to perturbations of the initial conditions has been observed.

Chapter 7

Conclusions and Outlook

This work is motivated by the need of developing a systematic methodology to improve the computational efficiency and physical interpretability of reduced-order models of chaotic systems. This is achieved by identifying the dominant subset of triadic interactions in Galerkin models of turbulent flows in such a way only a reduced number of terms is needed to temporally advance the system. The *a posteriori* and the *a priori* sparsification frameworks proposed here aim to tackle this problem from two different perspectives. The *a posteriori* technique follows a more classical approach from data, where the interactions are pruned based on their relative strength but the modal basis used to generate the reduced-order model is not modified. A possible drawback of such an approach is the loss between the modal structures and the energy fluxes due to the modification of the entries of $\mathbf{Q}$. To fill this conceptual gap the *a priori* sparsification approach was proposed. It leverages the key idea of encoding the information about the sparsity of the triadic interactions directly in the modal basis, such that the system of ODEs obtained by Galerkin projection has already a sparse structure of the nonlinear interactions tensor. Finally, it is worth citing that a third sparsification approach, called greedy approach, was proposed. The underlying idea is to select the quadratic interaction based on their energetic relevance without the need of any *a posteriori* or *a priori* optimisation procedure. However, these models inherit all the temporal stability issues of the dense models and cannot be used for practical purposes, therefore, this approach has been discarded. In the following sections, the main similarities and differences between the *a priori* and the *a posteriori* sparsification methods are presented discussing possible outlooks and ideas for future research are discussed.

7.1 Similarities

Both approaches have been proved able to recover the correct physics of the energy interactions. More specifically, the *a posteriori* sparsification has been shown to identify a relevant subset of triadic interactions coherent with the dominant subset of triadic interactions identified by energy analysis. More interestingly, we observed that the relevant subset of triadic interactions are successfully identified if the well-posedness of the regression problem is guaranteed, regardless of the modal basis used to perform Galerkin projection. Further analysis of the structure of the sparse systems shows that the relevant energy pattern is correctly identified by both approaches (see 5.20,6.8 and 5.22, 6.10) and energy conservation properties of the non linear interactions are preserved. Interestingly, these properties are conserved in the sparse system without the need for the introduction of specific constraints. Arguably, this information is already contained in the dynamics of the flow introduced in the sparsification by the presence of the modal acceleration term. On the other hand, it has been observed that for the *a priori* sparsification the presence of the stability constraint is needed to obtain sparse systems identifying the right pattern of triadic interactions and predicting the correct temporal evolution. Generally, it has been observed that both sparsification approaches perform better in case of larger systems (corresponding to a larger amount of reconstructed turbulent kinetic energy) that can be sparsified more aggressively with less impact both on the reconstruction error and on the ability of the new basis functions of reconstructing kinetic energy. More specifically, for the *a posteriori* sparsification more sparse systems are obtained either when more resolved systems (larger number of modes) or higher Reynolds numbers are considered. For the *a priori* method better sparsification performances can be obtained by increasing the size of the model N, reducing the slope of the sparsification curves in the $\rho - e_N$, or the size of subspace within the rotation is performed M, affecting the minimal value of the density that can be obtained. Subsequently, the temporal integration performances have been evaluated. As expected, both approaches outperform the dense systems when it comes to the prediction of fluctuating kinetic energy. This is strictly true for time spans of a length comparable with the training dataset. This improvement is due to the training phase targeting the dynamics of the DNS simulation and the presence of the stability constraint (4.35b) for the *a posteriori* and the *a priori* approaches, respectively. This produces a reconstructed flow field consistent with the results obtained by the numeric simulations. More specifically, the optimisation procedure implicit in both approaches eliminates the nonphysical oscillations associated with an excess of predicted turbulent kinetic energy. Overall, the sparse systems have been shown to provide quite accurate predictions of the reconstructed flow field for all the systems at the Reynolds numbers considered. These are interesting results showing that the implicit tuning present in both approaches can balance both the unresolved physics associated with the smaller scales and the non-modelled pressure term for open flows.

7.2 Differences

Although these two approaches generate a sparsification pattern that is coherent with the energy analysis, some differences between these two approaches remain. First, comparing the sparsification curves obtained for the lid-driven cavity flow (common to both approaches), it can be observed that different values of optimal densities are reached by the two different techniques. Namely, in the *a priori* formulation we manage to reach values of the density as low as $\rho \sim 0.8$. Conversely, in the *a posteriori* framework value of the as low as $\rho \sim 0.3$. In addition, figure 6.3-(b) to (d) shows that it is not possible to sparsify the system without affecting, even marginally, the ability of the new set of modes of reconstructing energetic kinetic energy. This behaviour is not avoidable and is due to theoretical considerations about the energetic optimality of the POD. On the other hand, in the *a posteriori* framework, it seems possible to prune interactions from the system without affecting neither the reconstruction error ϵ nor the fluctuating amplitude of the temporal modes $a_i(t)$. Consequently, the value of the eigenvalues λ_i, proportional to the mean fluctuation amplitude of $a_i(t)$, is well preserved by the sparsification procedure. It is arguable that this difference is due to the fact that the *a priori* approach has a stronger tie back to the physics of the flow.

A more important difference between the two methodologies concerns the choice of the optimally sparse reduced-order model. In fact, for the *a posteriori* sparsification approach the regularisation weight can be chosen to be arbitrarily large resulting in a system where all the coefficients have been pruned ($\rho = 0$) with no physical meaning nor practical utility. The choice between all the possible systems obtained with different values of the regularisation weight needs a heuristic criterion to identify a system with a good trade-off between sparsity, accuracy and physical reliability. As an example, in this work, we proposed a methodology leveraging the idea of looking for the inflection point of the l_1 based pseudo density. Obviously, other selection criteria are possible. This highlights the problem of interpretability in classical supervised machine learning, where the performance of a given model can be quantified only after temporal integration. On the other hand, for the *a priori* sparsification technique, due to its mathematical formulation, the optimally sparse system can be identified as the system with the lowest density that still satisfies the sparsity and stability constraints in the optimisation problem. Since all the sparse systems are temporally stable, this allows to automatically select the optimally sparse system without any need of user intervention and, in the author opinion, is an interesting feature to develop a completely user-independent sparsification framework.

A crucial difference between the two approaches concerns the long term temporal stability of the sparse systems and their robustness with respect to perturbations of the initial conditions. More specifically, it has been observed that the *a priori* sparsified systems have good long term temporal stability and, at the same time, good robustness with respect to the perturbation of the initial condition. Conversely, the systems obtained with the *a posteriori* approach seems to produce accurate results only for initial conditions select inside the dataset used to generate the reduced-order model and time intervals comparable with the length of the training dataset. In fact, for time intervals longer than the training dataset a progressive drift, away from the correct value o turbulent kinetic energy is observed. This behaviour shows the inception of time instability restricting the applicability of the sparse system for practical applications.

Lastly, few considerations about the computational complexity of the two algorithms. One of the key aspects of the *a posteriori* sparsification consists in recasting the heuristic problem of selecting the most relevant interactions in Galerkin models into a convex optimisation problem. Crucially, this problem is characterised by a unique minimum that can be efficiently found with gradient descent-based techniques. More specifically, for a model obtained by N modes the computational complexity of solving is $\mathcal{O}(N^3)$. This complexity makes the solution of the problem feasible for models of the size of few hundred of modes. Conversely, the formulation of the *a priori* sparsification framework requires a higher computational effort. More specifically, an analysis of the operation involved shows that only the evaluation of the derivative of the gradient requires $\mathcal{O}(M^3N)$ leading to a much higher computational cost with respect to its *a posteriori* counterpart. In the present work, we have been able to sparsify systems up to fifty modes rotated in a subspace of dimension $M = 100$. Lastly, while the *a posteriori* regression-based approach is based on well known and tested, the *a priori* approach is, to the author knowledge, the first attempt to develop such a framework. Therefore it is reasonable to believe that there is still room for improvement to scale this framework up to computational performances comparable with regression based techniques.

7.3 Application Guidelines

Overall, both methodologies lead to similar results. We provide here some guidelines for the application of these two approaches in different situations. Analysing the *a posteriori* approach first, we feel that this could be a good solution when Galerkin systems constructed by a large number of modes N are considered and the cost of the solution of the optimisation problem is relevant. In addition, the *a posteriori* sparsified systems are reliable only on time integration spans comparable with the length of the training dataset. Therefore, an example of a good fit for the application of the *a posteriori* sparsification is the online computation of the sparse model from streaming data to predict the behaviour of the system for future times. For this application, it is of paramount importance to have a sparsification algorithm that is computationally efficient to compute the system coefficients faster than the advancement of the CFD solution. Finally, the *a posteriori* sparsification framework could be a good solution when Galerkin systems are obtained by experimental data and only the values of the temporal modes $a_i(t)$ and their time derivative $\dot{a}_i(t)$ are easily obtainable while obtaining a good approximation of the matrix coefficients C_i,L_{ij} and Q_{ijk} is challenging.

Conversely, the *a priori* approach is computationally more expensive and requires the knowledge of all the components of the Galerkin system (4.6). At the same time, it provides a deeper insight into the physics and generates reduced-order models that are more reliable computationally. Therefore, we feel to suggest the use of this approach on flows obtained by numerical simulations where the coefficients matrices are easily obtainable. However, a more complex initial set up results in better performances making this approach advisable when it is important to obtain reduced-order model stable over long time spans. Lastly, one more consideration about the applications of these methodologies. One of the main targets of the initial phase of this work was to analyse the locality of the interaction in three-dimensional turbulence via sparse regression. However, it has been observed that the amount of data (more specifically $\dot{a}_i$) to obtain a reliable regression make the application of the *a posteriori* methodology not practical. We suggest that the *a priori* could be a possible solution to study high complex systems since only the amount of data needed to obtain a statistically converged modal decomposition is required.

7.4 Outlook

In this work, we provided a comprehensive description of the application of statistical learning techniques to reduced-order models of turbulent flows. Interestingly, it has been shown that different sparsification approaches can recover the correct physics of the triadic interactions for different flow configurations. Although these techniques seem to be promising, some avenues still need to be explored to obtain a more complete understanding and, eventually, apply the sparsification techniques in real-life flow applications.

From a computational point of view the *a posteriori* sparsification technique is quite well understood. The mathematical framework of the l_1 based regression is well tested and computationally efficient algorithms for its solution are available in the vast majority of machine learning libraries. Conversely, regarding the *a priori* approach several aspects deserve further investigation. More specifically, unlike in l_1-regression-based sparsification methods where the model coefficients are the optimisation variables, here these coefficients are cubic polynomial functions of the optimisation variables, the entries of the rotation matrix. The first consequence is that the optimisation problem is non-convex and many local optima, *i.e.* many different sets of basis functions, exist. It was observed that initial guesses close to the original POD basis, *i.e.* with good energy representation ability, repeatedly converge to the same optimal solution, which has a consistent physical interpretation. However, the lack of strong uniqueness guarantees (as for many other modal decomposition techniques) may render the approach questionable. Second, it was shown that it is not possible to shrink to zero an arbitrary number of coefficients since the link between the basis functions and the Galerkin model must always be maintained. As a result, the present *a priori* sparsification technique produces relatively denser Galerkin models than the *a posteriori* LASSO-based approach considered for the same test case. Rotating the original POD basis into higher-dimensional subspaces, with a ratio M/N higher than that considered in the present work, to further increase sparsity is possible, although it would necessarily result in increased computational costs for the optimisation.

From a more physical perspective, it has been shown that for a two-dimensional flow configuration the sparsification can generate reduced order models with computational complexity scaling as N^a with $a < 3$ instead of N^3, as for the dense models. The interaction between sparsification and the physics of the flow is certainly an aspect that deserves more investigation. More specifically, it would be of great interest to understand how the local nature of triadic interactions in three dimensional flows can be used to improve the computational efficiency of Galerkin models affecting only marginally their prediction capabilities. In addition, it is in the author opinion that the locality of the interaction in three-dimensional turbulence might be used advantageously to promote sparsification as the Reynolds number increases making high Reynolds number turbulent flows more prone to sparsification. Due to the complexity of three-dimensional turbulent flows, this aspect needs to be analysed in a dedicated study but, arguably, leveraging the flow physics to promote sparsification is the key to develop the next generation of high Reynolds number flows computational models.

From a more fundamental perspective, the meaning of looking for a sparse basis function by rotating the original basis function need to be discussed. In this work, it has been shown that a new basis functions generating Galerkin models with sparse matrices coefficients can be obtained by a small rotation of the original POD basis. However, it has also been observed that the sparsification is less effective than for the *a posteriori* approach. Crucially, the uniqueness of the resulting basis functions cannot be guaranteed if the initial guess in the optimisation problem is perturbed enough from the initial condition. In addition, all the different optimal solutions have been observed to be only marginally different. This result brings up the question of whether an approach based on a small modification of an originally dense basis function is the best approach to promote sparsity in the basis function. A different technique can be developed by the observation that sparsity of the energetic interactions in modal space results from flow structure that interacts locally in physical space. Therefore, a possible approach to generate a modal basis naturally incorporating the sparsity could be achieved by developing a methodology that splits up the domain in not-overlapping partitions then generate a basis function defined only inside these partitions and zero elsewhere. These locally defined basis functions can naturally incorporate the local nature of flow structure interactions.

Finally, looking into possible future applications, the author feels to suggest that this set of methodologies could serve as the key enabler for autonomous air and sea vehicles. Physics–informed models may be integrated into a physical models, for onboard and real–time applications. In this context, the physical model uses information derived from a network of sensors embedded in the system, at certain spatial locations. The measured information drives the reduced-order model which reconstructs the flow state around the system and feeds the state reconstruction to a mission manager. This type of application is feasible in the near future by integrating commercial off–the–shelves hardware components into a flying testbed.

Appendix A

A priori objective function formulation

This section is dedicated to the formulation of the objective function for the optimisation problem in the *a priori* sparsification technique.

A.1 Preliminaries

For sake of generality, we introduce the matrix $\mathbf{W}$ containing information about the mesh (if not uniform) and the space discretisation of the gradient and laplacian operators. The following relationships hold:

$$\mathbf{D} = \mathbf{U}\mathbf{A}^\top \tag{A.1}$$

$$\mathbf{I} = \mathbf{U}^\top\mathbf{W}\mathbf{U} = \mathbf{A}\mathbf{A}^\top \tag{A.2}$$

$$\tilde{\mathbf{U}} = \mathbf{U}\mathbf{X} \tag{A.3}$$

$$\tilde{\mathbf{A}} = \mathbf{A}\mathbf{X} \tag{A.4}$$

$$\mathbf{C} = \mathbf{D}^\top\mathbf{W}\mathbf{D} = \mathbf{A}\mathbf{U}^\top\mathbf{W}\mathbf{U}\mathbf{A}^\top = \mathbf{A}\mathbf{A}^\top \tag{A.5}$$

$$\tilde{\mathbf{C}} = \tilde{\mathbf{D}}^\top\mathbf{W}\tilde{\mathbf{D}} = \tilde{\mathbf{A}}\tilde{\mathbf{U}}^\top\mathbf{W}\tilde{\mathbf{U}}\tilde{\mathbf{A}}^\top = \tilde{\mathbf{A}}\mathbf{X}^\top\underbrace{\mathbf{U}^\top\mathbf{W}\mathbf{U}}_{\mathbf{I}}\mathbf{X}\tilde{\mathbf{A}}^\top = \tilde{\mathbf{A}}\tilde{\mathbf{A}}^\top \tag{A.6}$$

A.2 Objective Function Formulation

Recalling that all matrices in the rotated basis are defined with a tilde, we define the error of the new basis function with respect to the original dataset as $\mathbf{E} = \mathbf{D} - \tilde{\mathbf{D}}$ we want to express the mean quadratic error as a function of **X**. Let's consider:

$$Obj = \frac{1}{N_T} Tr(\mathbf{E}^\top \mathbf{W} \mathbf{E}) = \frac{1}{N_T} Tr((\mathbf{D} - \tilde{\mathbf{D}})^\top \mathbf{W} (\mathbf{D} - \tilde{\mathbf{D}})), \tag{A.7}$$

that can be expanded as follows

$$Obj = \frac{1}{N_T} Tr(\underbrace{\mathbf{D}^\top \mathbf{W} \mathbf{D}}_{I} + \underbrace{\tilde{\mathbf{D}}^\top \mathbf{W} \tilde{\mathbf{D}}}_{II} - 2 \underbrace{\mathbf{D}^\top \mathbf{W} \tilde{\mathbf{D}}}_{III}). \tag{A.8}$$

We break down this expression into three pieces: I,II and III as,

$$I : \mathbf{D}^\top \mathbf{W} \mathbf{D} = \mathbf{A} \mathbf{U}^\top \mathbf{W} \mathbf{U} \mathbf{A}^\top = \mathbf{A} \mathbf{A}^\top, \tag{A.9}$$

$$II : \tilde{\mathbf{D}}^\top \mathbf{W} \tilde{\mathbf{D}} = \tilde{\mathbf{A}} \tilde{\mathbf{U}}^\top \mathbf{W} \tilde{\mathbf{U}} \tilde{\mathbf{A}}^\top = \mathbf{A} \mathbf{X} \mathbf{X}^\top \underbrace{\mathbf{U}^\top \mathbf{W} \mathbf{U}}_{\mathbf{I}} \mathbf{X} \mathbf{X}^\top \mathbf{A}^\top = \mathbf{A} \mathbf{X} \mathbf{X}^\top \mathbf{X} \mathbf{X}^\top \mathbf{A}^\top = \mathbf{A} \mathbf{X} \mathbf{X}^\top \mathbf{A}^\top, \tag{A.10}$$

$$III : \mathbf{D}^\top \mathbf{W} \tilde{\mathbf{D}} = \mathbf{A} \mathbf{U}^\top \mathbf{W} \tilde{\mathbf{U}} \tilde{\mathbf{A}}^\top = \mathbf{A} \mathbf{U}^\top \mathbf{W} \mathbf{U} \mathbf{X} \mathbf{X}^\top \mathbf{A}^\top = \mathbf{A} \mathbf{X} \mathbf{X}^\top \mathbf{A}^\top. \tag{A.11}$$

Exploiting the property of the trace operator of being invariant under circular permutation of its factors we obtain,

$$Tr(\mathbf{D}^\top \mathbf{W} \mathbf{D}) = Tr(\mathbf{A} \mathbf{A}^\top) = Tr(\mathbf{A}^\top \mathbf{A}) = Tr(N_T \mathbf{\Lambda}), \tag{A.12}$$

$$Tr(\tilde{\mathbf{D}}^\top \mathbf{W} \tilde{\mathbf{D}})) = Tr(\mathbf{X}^\top \mathbf{A}^\top \mathbf{A} \mathbf{X}), \tag{A.13}$$

$$Tr(\mathbf{D}^\top \mathbf{W} \tilde{\mathbf{D}}) = Tr(\mathbf{X}^\top \mathbf{A}^\top \mathbf{A} \mathbf{X}). \tag{A.14}$$

Therefore the objective function reads

$$Obj = \frac{1}{N_T} (Tr(\mathbf{A}^\top \mathbf{A}) - Tr(\mathbf{X}^\top \mathbf{A}^\top \mathbf{A} \mathbf{X})), \tag{A.15}$$

that can be simplified as

$$Obj = \frac{1}{N_T} (Tr(\mathbf{A}^\top \mathbf{A}) - Tr(\mathbf{X}^\top \mathbf{A}^\top \mathbf{A} \mathbf{X})) = Tr(\mathbf{\Lambda}) - Tr(\mathbf{X}^\top \mathbf{\Lambda} \mathbf{X})) = Tr(\mathbf{\Lambda}) - Tr(\tilde{\mathbf{\Lambda}}). \tag{A.16}$$

Appendix B

Gradients evaluation

B.1 Gradient of the Objective Function

Recalling the definition of the objective function:

$$Obj = Tr(\mathbf{\Lambda} - \mathbf{X}^\top \mathbf{\Lambda} \mathbf{X}) \tag{B.1}$$

Since the eigenvalues matrix is diagonal and $\mathbf{\Lambda}^\top = \mathbf{\Lambda}$ holds. Its gradient can be easily computed as:

$$\frac{\partial Obj}{\partial \mathbf{X}} = -(\mathbf{\Lambda}^\top + \mathbf{\Lambda})\mathbf{X} = -2\mathbf{\Lambda}^\top \mathbf{X} \tag{B.2}$$

B.2 Gradient of the stability constraint

The temporal stability constraint is defined as $f(\mathbf{X}) = 0$ where

$$f(\mathbf{X}) = Tr(\mathbf{X}^\top \mathbf{L} \mathbf{X}) - \eta \tag{B.3}$$

In the present work the constant η is defined as $\eta = -|tr(\mathbf{L})|\eta_0$, with η_0 having increasing values starting from 1.

$$\frac{\partial f(\mathbf{X})}{\partial \mathbf{X}} = \frac{\partial}{\partial \mathbf{X}} Tr(\mathbf{X}^\top \mathbf{L} \mathbf{X}) = \frac{\partial}{\partial \mathbf{X}} (\mathbf{X}^\top \mathbf{L} \mathbf{X}) = (\mathbf{L}^\top + \mathbf{L})\mathbf{X} \tag{B.4}$$

B.3 l_1 part of the objective function

The sparsity promoting constraint is defined with the following inequality

$$g(\mathbf{X}) = ||\tilde{\mathbf{Q}}||_1 - c \leq 0 \tag{B.5}$$

The tensor $\tilde{\mathbf{Q}}$ is computed as a rotation of the tensor $\mathbf{Q}$ as in 4.27. For the sake of interpretability in the present work we defined the constant c as $c = ||\mathbf{Q}||_1/\xi$.

The derivative of the constrain with respect the rotation matrix $\mathbf{X}$ is defined as follows

$$\frac{\partial g(\mathbf{X})}{\partial \mathbf{X}} = \frac{\partial ||\tilde{Q}_{ijk}||_1}{\partial X_{mn}} = \frac{\partial \tilde{Q}_{ijk}}{\partial X_{mn}} \frac{\tilde{Q}_{ijk}}{|\tilde{Q}_{ijk}|} \tag{B.6}$$

to compute the derivative of $\tilde{Q}_{ijk}$ we proceed computing its variation with respect to a small variation of $\mathbf{X}$ defined as $\mathbf{X} + d\mathbf{X} = \mathbf{X} + \epsilon\mathbf{X}$

$$\frac{\partial \tilde{\mathbf{Q}}}{\partial \mathbf{X}} = \lim_{\epsilon \to 0} \frac{\tilde{\mathbf{Q}}(\mathbf{X} + \epsilon\mathbf{X}) - \tilde{\mathbf{Q}}(\mathbf{X})}{\epsilon}, \tag{B.7}$$

expanding out this expression in Einstein notation and retaining only the first order terms in ϵ we obtain:

$$d\tilde{Q}_{ijk} = Q_{pqr}(X_{pi}X_{qj}dX_{kr} + X_{pi}X_{rk}dX_{jq} + X_{rk}X_{qj}dX_{pi}). \tag{B.8}$$

Exploiting the properties of the Kronecker's delta $dX_{kr} = dX_{mn}\delta_{km}\delta_{rm}$ to rewrite the variations with respect the indexes m, n, expression (B.8) can be rewritten as:

$$d\tilde{Q}_{ijk} = \frac{\partial \tilde{Q}_{ijk}}{\partial X_{mn}} dX_{mn} = Q_{pqr}(X_{pi}X_{qj}\delta_{km}\delta_{rm} + X_{pi}X_{rk}\delta_{jn}\delta_{qm} + X_{rk}X_{qj}\delta_{in}\delta_{pm})dX_{mn}, \tag{B.9}$$

that after some algebra becomes,

$$\frac{\partial \tilde{Q}_{ijk}}{\partial X_{mn}} = Q_{pqr}(X_{pi}X_{qj}\delta_{kn}\delta_{rm} + X_{pi}X_{rk}\delta_{jn}\delta_{qm} + X_{rk}X_{qj}\delta_{in}\delta_{pm}) \tag{B.10}$$

Plugging this expression in (B.3) we obtain the analytical expression of the gradient of the sparsity promoting constraint. Crucially, in this formulation of the *a priori* sparsification all the gradients are known and easily evaluating analytically enabling good scalability of the *a priori* algorithm to reduced-order models of medium to large sizes.

Appendix C

Equivalence between momentum and energy LASSO

In this section, it is shown how to rewrite the LASSO problem in terms of the energy equation and we will prove that the solutions of the lasso in terms of the energy equation and of the momentum equation are equivalent. First, we recall the definition of the momentum and energy equation in modal space.

$$\dot{a}_i(t) = \mathsf{C}_i + \sum_{j=1}^{N} \mathsf{L}_{ij} a_j(t) + \sum_{j=1}^{N}\sum_{k=1}^{N} \mathsf{Q}_{ijk} a_j(t) a_k(t) \qquad i = 1\ldots, N, \tag{C.1}$$

and

$$\dot{e}_i(t) = \mathsf{C}_i a_i(t) + \sum_{j=1}^{N} \mathsf{L}_{ij} a_i(t) a_j(t) + \sum_{j=1}^{N}\sum_{k=1}^{N} \mathsf{Q}_{ijk} a_i(t) a_j(t) a_k(t), \qquad i = 1, \ldots, N, \tag{C.2}$$

with $\dot{e}_i(t) = a_i(t)\dot{a}_i(t)$.

We have already shown that (C.1) can be rewritten in matrix form defining the modal acceleration matrix $\dot{\mathbf{A}}$ (whose column is defined as $\dot{\mathbf{A}}_\mathbf{i}$) and the database matrix $\mathbf{\Theta}(\mathbf{A})$. To rewrite (C.2) in a way formally identical to (4.22) we defined the column-wise multiplication operator $\boxtimes$ that given a vector $\mathbf{v} = [v_1 \quad v_2 \quad v_3]$ and a matrix $\mathbf{M} \in \Re^{3\times 3}$ defined as:

$$\mathbf{v} \boxtimes \mathbf{M} = \begin{pmatrix} v_1 & 0 & 0 \\ 0 & v_2 & 0 \\ 0 & 0 & v_3 \end{pmatrix} \begin{pmatrix} m_{11} & m_{12} & m_{13} \\ m_{21} & m_{22} & m_{23} \\ m_{31} & m_{32} & m_{33} \end{pmatrix} = \begin{pmatrix} v_1 a_{11} & v_1 a_{12} & v_1 a_{13} \\ v_2 a_{21} & v_2 a_{22} & v_2 a_{23} \\ v_3 a_{31} & v_3 a_{32} & v_3 a_{33} \end{pmatrix}. \tag{C.3}$$

In the present case, the vector $\mathbf{v}$ is the modal oscillating amplitude $\mathbf{A}_i$ while the matrix $\mathbf{M}$ is the modal acceleration matrix $\dot{\mathbf{A}}_i$, for the right-hand side, and the database matrix $\mathbf{\Theta}(\mathbf{A})$, for the left-hand side, respectively. The energy equation can be rewritten in matrix form for each mode exploiting the properties of the operator $\boxtimes$ and the matrices $\dot{\mathbf{A}}_i$ and $\boldsymbol{\Theta}(\mathbf{A})$.

$$\mathbf{A}_i \boxtimes (\dot{\mathbf{A}}_{\mathbf{i}} - \boldsymbol{\Theta}(\mathbf{A})\boldsymbol{\beta_i}) = 0, \qquad i = 1, \dots, N. \tag{C.4}$$

Therefore the LASSO problem can be formulated as

$$\min_{\boldsymbol{\beta}_i} ||\boldsymbol{\Theta}_E(\mathbf{A})\boldsymbol{\beta}_i - \dot{\mathbf{E}}_i||_2^2 + \gamma_i ||\boldsymbol{\beta}_i||_1, \qquad i = 1, ..., N, \tag{C.5}$$

with $\dot{\mathbf{E}}_i = \mathbf{A}_i \boxtimes \dot{\mathbf{A}}_{\mathbf{i}}$ and $\boldsymbol{\Theta}_E = \mathbf{A}_i \boxtimes \boldsymbol{\Theta}$.

Since the solution of the LASSO can be computed as iteratively thresholded least-squares we can compute analytically the solution of the least square part as:

$$\frac{\partial}{\partial \boldsymbol{\beta}_i} ||\boldsymbol{\Theta}_E(\mathbf{A})\boldsymbol{\beta}_i - \dot{\mathbf{E}}_i||_2^2 = 0 \tag{C.6}$$

since the $\boxtimes$ satisfies the distributive property and it is orthogonal with respect to the $\partial / \partial \boldsymbol{\beta}_i$ operator, the following relation holds

$$||\mathbf{A}_i||_2^2 \boxtimes \frac{\partial}{\partial \boldsymbol{\beta_i}} ||\boldsymbol{\Theta}(A)\boldsymbol{\beta}_i - \dot{A}_i||_2^2 = 0. \tag{C.7}$$

The former expression has solution $\boldsymbol{\beta}_i = (\Theta^\top \Theta)^{-1} \Theta^\top \dot{\mathbf{A}}_i$. Therefore performing LASSO regression and on the energy equation bring to the same solution for $\boldsymbol{\beta}_i$. Geometrically, it can be shown that, the solution to the LASSO problem for the energy equation is equivalent to the one of the momentum equation, since the value of the gradient is scaled mode by mode by the modal amplitude $||\mathbf{A}_i||_2^2$ but the location of the minima in $\boldsymbol{\beta}_i$ space is preserved.

Appendix D

About the definition of the sparsity promoting constraint

An aspect of paramount importance, in the development of the *a priori* sparsification algorithm, is the definition of the sparsity promoting constraint (4.35b). The philosophy of this approach is to obtain a sparse approximation of the tensor **Q** exploiting the idea that a set of dominant energy interactions exists in **N**. To this goal, in the formulation presented in this work, an optimisation problem with an inequality constraint on the l_1 norm of the rotated quadratic interaction tensor $\tilde{\mathbf{Q}}$ is solved. However, other approaches are possible and, in a preliminary phase of this work, a formulation based on the penalisation of the rotated triadic interaction tensor $\tilde{N}_{ijk} = \tilde{Q}_{ijk}\overline{\tilde{a}_i\tilde{a}_j\tilde{a}_k}$ was tested. The formulation of the new problem is totally analogous to the formulation (4.30) with the only difference of a sparsity promoting constraint defined with respect to the triadic interaction tensor **N** instead of **Q** as

$$||\tilde{\mathbf{N}}||_1 \leq \zeta^{-1}||\mathbf{N}||_1 \tag{D.1}$$

Analogously, to what was done before to understand the role of this constraint in the solution of problem (4.35) we consider a small reduced-order model of the lid-driven cavity flow at $Re = 2 \times 10^4$. Thus we visualise the effects of the rotation varying two entries of **X** without solving the optimisation problem (4.35). Panel (a) in figure D.1 shows the contours of $||\tilde{\mathbf{N}}||_1/||\mathbf{N}||_1$. Similarly to what observed in figure 6.1 the constraint displays several valley of local minima arising from non-convex, non-smooth nature of the quantity $||\tilde{\mathbf{N}}||_1$. To better visualise this feature in panel (b) two horizontal and one vertical cut (corresponding to the lines drawn in panel (a)) are displayed. As observed previously in figure 6.1-(b,c) each minimum has a discontinuous first derivative corresponding to one of more entries of $\tilde{\mathbf{N}}_{ijk}$ changing sign. The general idea is completely analogous to the one illustrated in figure 6.1 for the sparsification of $\tilde{\mathbf{Q}}$. The

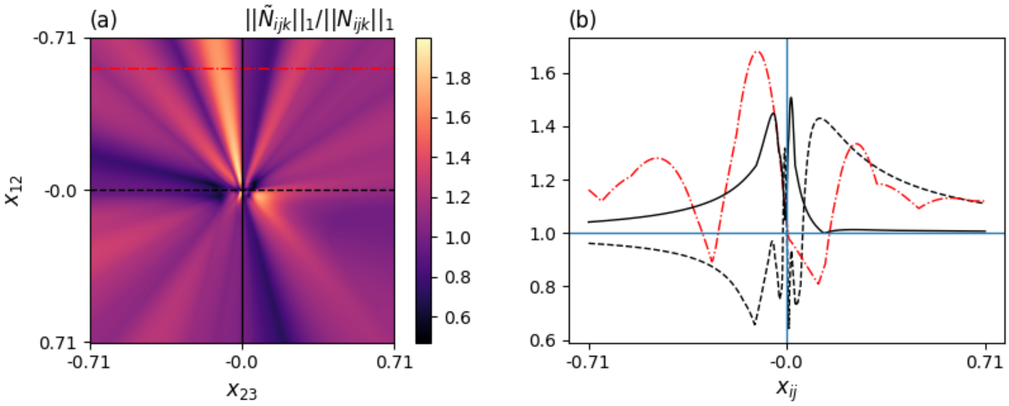

FIGURE D.1: Panel (a): map of $||\tilde{N}_{ijk}||_1/||N_{ijk}||_1$ projected on the $X_{12} - X_{23}$ plane. Panel (b): Value of $||\tilde{N}_{ijk}||_1/||N_{ijk}||_1$ along three different cuts corresponding to the line drawn in panel (a).

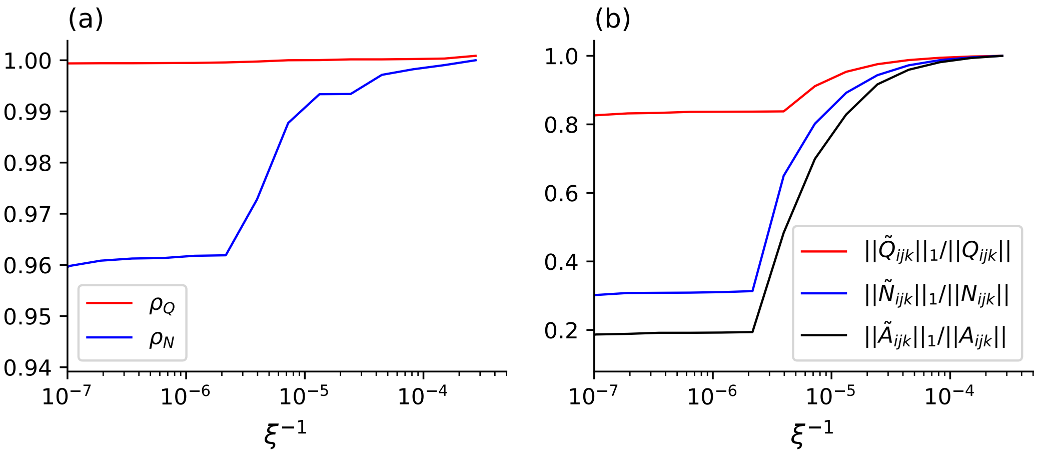

FIGURE D.2: Panel (a): value of the density defined with respect to the tensor Q and the tensor N against the inverse of the regularisation weight ξ^{-1}. Panel (b): Value of the l_1 norm for $\tilde{Q}_{ijk}$,$\tilde{N}_{ijk}$ and $\tilde{A}_{ijk} = \overline{\tilde{a}_i\tilde{a}_j\tilde{a}_k}$.

increase of the penalisation weight ξ^{-1} pushes $||\tilde{\mathbf{N}}||_1$ closer to its minima and consequently one or more of its entries $\tilde{N}_{ijk}$ are shrunk to zero due to the non-differentiable nature of the l_1 norm. From the computational point of view all the considerations done in section 4.3.2 are still valid, since the triadic interaction tensor can be written as element to element multiplication of the two tensors $\tilde{Q}_{ijk}$ and $\tilde{A}_{ijk} = \overline{\tilde{a}_i\tilde{a}_j\tilde{a}_k}$. The gradient is the derivative of a product and can be computed as

$$\frac{\partial \tilde{N}_{ijk}}{\partial X_{mn}} = \frac{\partial \tilde{Q}_{ijk}}{\partial X_{mn}}\overline{\tilde{a}_i\tilde{a}_j\tilde{a}_k} + \tilde{Q}_{ijk}\frac{\partial \overline{\tilde{a}_i\tilde{a}_j\tilde{a}_k}}{\partial X_{mn}}, \tag{D.2}$$

Consequently, all the operations considered in section 4.3.2 need to be performed twice since two different third order tensors are involved. Thus, the scaling in the function of M, N of the computational cost is analogous to the one discussed previously for the definition of the sparsity constraint in terms of $\mathbf{Q}$. To better understand how the introduction of this sparsity promoting constraint affects the solution of (4.35), we solved the problem with the constraint (D.1) with $N = 30$, $M/N = 2$ and increasing values of the penalisation weight ξ. Figure D.2-(a) shows the density ρ as defined in (4.36) of the tensors $\mathbf{Q}$ and $\mathbf{N}$ called ρ_Q and ρ_N, respectively. It can be observed that only the tensor $\tilde{\mathbf{N}}$ is effectively sparsified while the density of $\tilde{\mathbf{Q}}$ does not vary appreciably while increasing the penalisation ξ^{-1}. Figure D.2-(b) shows the l_1 norm of $\tilde{\mathbf{Q}}$,$\tilde{\mathbf{N}}$ and $\tilde{A}_{ijk} = \overline{\tilde{a}_i\tilde{a}_j\tilde{a}_k}$ normalised with respect to the original value to visualise the decrease in percentage with respect the starting point $\mathbf{X}_0 = \mathbf{I}$. Interestingly, we observe that the optimisation decreases the value of $\tilde{\mathbf{N}}$ affects, in relative terms, considerably $\tilde{\mathbf{A}}$ and only marginally $\tilde{\mathbf{Q}}$. Arguably, this is the reason why $\mathbf{N}$ is much more effectively sparsified with respect to $\mathbf{Q}$ when the optimisation problem is solved. In the light of these results, it has been chosen not to pursue this formulation since we are aiming to generate a reduced order model with a sparse structure of the $\mathbf{Q}$. A sparsification of only the tensor N_{ijk} would not produce any gain neither for the interpretability of the system nor in computational costs. Nevertheless, this example shows the flexibility of the subspace rotation approach and how it can be used to enforce/extract physical features from the mathematical structure of reduced-order models.

References

Frédéric Alizard and Jean-Christophe Robinet. Modeling of optimal perturbations in flat plate boundary layer using global modes: benefits and limits. *Theoretical and Computational Fluid Dynamics*, 25(1-4):147–165, 2011.

David Amsallem and Charbel Farhat. Interpolation method for adapting reduced-order models and application to aeroelasticity. *AIAA journal*, 46(7):1803–1813, 2008.

David Amsallem and Charbel Farhat. Stabilization of projection-based reduced-order models. *International Journal for Numerical Methods in Engineering*, 91(4):358–377, 2012.

H. Arbabi and I. Mezić. Study of dynamics in post-transient flows using Koopman mode decomposition. *Physical Review Fluids*, 2(12):124402, 2017.

F. Auteri, N. Parolini, and L. Quartapelle. Numerical investigation on the stability of singular driven cavity flow. *Journal of Computational Physics*, 183(1):1–25, 2002.

M. Balajewicz, E. Dowell, and B. Noack. Low-dimensional modelling of high-Reynolds-number shear flows incorporating constraints from the Navier–Stokes equation. *Journal of Fluid Mechanics*, 729, 2013.

Maciej Balajewicz, Irina Tezaur, and Earl Dowell. Minimal subspace rotation on the stiefel manifold for stabilization and enhancement of projection-based reduced order models for the compressible navier–stokes equations. *Journal of Computational Physics*, 321:224–241, 2016.

M Bergmann, C-H Bruneau, and A Iollo. Improvement of reduced order modeling based on pod. In *Computational Fluid Dynamics 2008*, pages 779–784. Springer, 2009.

Michel Bergmann, Charles-Henri Bruneau, and Angelo Iollo. Improvement of Reduced Order Modeling based on Proper Orthogonal Decomposition. Research Report RR-6561, INRIA, 2008. URL `https://hal.inria.fr/inria-00288415`.

A. L. Blum and P. Langley. Selection of relevant features and examples in machine learning. *Artificial intelligence*, 97(1-2):245–271, 1997.

V. B. L. Boppana and J. S. B. Gajjar. Global flow instability in a lid-driven cavity. *International Journal for Numerical Methods in Fluids*, 62(8):827–853, 2010.

O. Botella and R. Peyret. Benchmark spectral results on the lid-driven cavity flow. *Computers & Fluids*, 27(4):421 – 433, 1998.

J-L Bourguignon, JA Tropp, AS Sharma, and BJ McKeon. Compact representation of wall-bounded turbulence using compressive sampling. *Physics of Fluids*, 26(1): 015109, 2014.

J. Brasseur and C. Wei. Interscale dynamics and local isotropy in high reynolds number turbulence within triadic interactions. *Physics of Fluids*, 6, 1994.

S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15), 2016.

S. L. Brunton, B. R. Noack, and P. Koumoutsakos. Machine learning for fluid mechanics. *Annual Review of Fluid Mechanics*, 52, 2019.

Tan Bui-Thanh, Karen Willcox, and Omar Ghattas. Parametric reduced-order models for probabilistic analysis of unsteady aerodynamic applications. *AIAA journal*, 46 (10):2520–2529, 2008.

Alfonso Caiazzo, Traian Iliescu, Volker John, and Swetlana Schyschlowa. A numerical investigation of velocity–pressure reduced order models for incompressible flows. *Journal of Computational Physics*, 259:598–616, 2014.

W. Cazemier, R. W. C. P. Verstappen, and A. E. P. Veldman. Proper orthogonal decomposition and low-dimensional models for driven cavity flows. *Physics of fluids*, 10(7): 1685–1699, 1998.

K. K. Chen, J. H. Tu, and C. W. Rowley. Variants of Dynamic Mode Decomposition: Boundary Condition, Koopman, and Fourier Analyses. *Journal of Nonlinear Sciences*, 22(6):887–915, April 2012.

L. Cordier, B. A. El Majd, and J. Favier. Calibration of POD reduced-order models using Tikhonov regularization. *International Journal for Numerical Methods in Fluids*, 63(2): 269–296, 2010.

M. Couplet, P. Sagaut, and C. Basdevant. Intermodal energy transfers in a proper orthogonal decomposition–galerkin representation of a turbulent separated flow. *Journal of Fluid Mechanics*, 491, 2003.

M. Couplet, C. Basdevant, and P. Sagaut. Calibrated reduced-order POD-Galerkin system for fluid flow modelling. *Journal of Computational Physics*, 207(1):192–220, 2005.

S. Diamond and S. Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83), 2016.

J Andrzej Domaradzki and Robert S Rogallo. Local energy transfer and nonlocal interactions in homogeneous, isotropic turbulence. *Physics of Fluids A: Fluid Dynamics*, 2 (3):413–426, 1990.

JA Domaradzki and Nikolaus A Adams. Direct modelling of subgrid scales of turbulence in large eddy simulations. *Journal of Turbulence*, 3(024):1, 2002.

L. Fick, Y. Maday, A. T. Patera, and T. Taddei. A stabilized POD model for turbulent flows over a range of reynolds numbers: Optimal parameter sampling and constrained projection. *Journal of Computational Physics*, 371:214 – 243, 2018.

J. Friedman, T. Hastie, and R. J. Tibshirani. *The Elements of Statistical Learning*. Springer, 2008.

B Galletti, CH Bruneau, Luca Zannetti, and Angelo Iollo. Low-order modelling of laminar flow regimes past a confined square cylinder. *Journal of Fluid Mechanics*, 503: 161–170, 2004.

F Gómez, HM Blackburn, M Rudman, BJ McKeon, M Luhar, R Moarref, and AS Sharma. On the origin of frequency sparsity in direct numerical simulations of turbulent pipe flow. *Physics of Fluids*, 26(10):101703, 2014.

F Gómez, HM Blackburn, M Rudman, AS Sharma, and BJ McKeon. A reduced-order model of three-dimensional unsteady flow in a cavity based on the resolvent operator. *Journal of Fluid Mechanics*, 798, 2016.

Per Christian Hansen. The discrete picard condition for discrete ill-posed problems. *BIT Numerical Mathematics*, 30(4):658–672, 1990.

T. Hastie, R. Tibshirani, and M. Wainwright. *Statistical learning with sparsity: the LASSO and generalizations*. Chapman and Hall/CRC, 2015.

Philip J Holmes, John L Lumley, Gal Berkooz, Jonathan C Mattingly, and Ralf W Wittenberg. Low-dimensional models of coherent structures in turbulence. *Physics Reports*, 287(4):337–384, 1997.

J. Jiménez. Coherent structures in wall-bounded turbulence. *Journal of Fluid Mechanics*, 842, 2018.

Steven G Johnson. The nlopt nonlinear-optimization package, 2014.

M. R. Jovanović, P. J. Schmid, and J. W. Nichols. Low-rank and sparse Dynamic Mode Decomposition. *Center for Turbulence Research Annual Research Briefs*, 2012:139–152, 2012.

M. R. Jovanović, P. J. Schmid, and J. W. Nichols. Sparsity-promoting Dynamic Mode Decomposition. *Physics of Fluids*, 26(2):024103, 2014.

Vassili Kitsios, Laurent Cordier, J-P Bonnet, Andrew Ooi, and Julio Soria. On the coherent structures and stability properties of a leading-edge separated aerofoil with turbulent recirculation. *Journal of Fluid Mechanics*, 683:395–416, 2011.

R. H. Kraichnan. Inertial-range transfer in two-and three-dimensional turbulence. *Journal of Fluid Mechanics*, 47(3):525–535, 1971.

Toni Lassila, Alfio Quarteroni, and Gianluigi Rozza. A reduced basis model with parametric coupling for fluid-structure interaction problems. *SIAM Journal on Scientific Computing*, 34(2):A1187–A1213, 2012.

J. P. Laval, B. Dubrulle, and S. Nazarenko. Nonlocality of interaction of scales in the dynamics of 2D incompressible fluids. *Physical Review Letters*, 83(20):4061, 1999.

J. C. Loiseau and S. L. Brunton. Constrained sparse Galerkin regression. *Journal of Fluid Mechanics*, 838:42–67, 2018.

J. L. Lumley. Stochastic tools in turbulence. Technical report, 1970.

J. L. Lumley. Computational modeling of turbulent flows. In *Advances in applied mechanics*, volume 18, pages 123–176. Elsevier, 1979.

John L Lumley. *Stochastic tools in turbulence*. Courier Corporation, 2007.

M. A. Mendez, M. Balabane, and J.-M. Buchlin. Multi-scale proper orthogonal decomposition of complex fluid flows. *Journal of Fluid Mechanics*, 870:988–1036, 2019. .

I. Mezić. Analysis of fluid flows via spectral properties of the Koopman operator. *Annual Review of Fluid Mechanics*, 45:357–378, 2013.

R Moarref, MR Jovanović, JA Tropp, AS Sharma, and BJ McKeon. A low-order decomposition of turbulent channel flow via resolvent analysis and convex optimization. *Physics of Fluids*, 26(5):051701, 2014.

A. G. Nair and K. Taira. Network-theoretic approach to sparsified discrete vortex dynamics. *Journal of Fluid Mechanics*, 768:549–571, 2015.

Mark EJ Newman. Fast algorithm for detecting community structure in networks. *Physical review E*, 69(6):066133, 2004.

V.B. Nguyen, M. Buffoni, K. Willcox, and B.C. Khoo. Model reduction for reacting flow applications. *International Journal of Computational Fluid Dynamics*, 28(3-4):91–105, 2014. . URL `https://doi.org/10.1080/10618562.2014.911848`.

B. Noack, M. Morzynski, and G. Tadmor. *Reduced-order modelling for flow control*, volume 528. Springer Science & Business Media, 2011.

B. R. Noack and H. Eckelmann. A Global Stability Analysis of the Steady and Periodic Cylinder Wake. *Journal of Fluid Mechanics*, 270:297–330, January 1994.

B. R. Noack, W. Stankiewicz, M. Morzyński, and P. J. Schmid. Recursive dynamic mode decomposition of transient and post-transient wake flows. *Journal of Fluid Mechanics*, 809:843–872, November 2016.

Bernd R Noack, Paul Papas, and Peter A Monkewitz. The need for a pressure-term representation in empirical galerkin models of incompressible shear flows. *Journal of Fluid Mechanics*, 523:339–365, 2005.

Bernd R Noack, Michael Schlegel, Boye Ahlborn, Gerd Mutschke, Marek Morzyński, Pierre Comte, and Gilead Tadmor. A finite-time thermodynamics of unsteady fluid flows. *Journal of Non-Equilibrium Thermodynamics*, 33(2):103–148, 2008.

K. Ohkitani. Nonlocality in a forced two-dimensional turbulence. *Physics of Fluids A: Fluid Dynamics*, 2(9):1529–1531, 1990.

J. Östh, B. R. Noack, S. Krajnović, D. Barros, and J. Borée. On the need for a nonlinear subscale turbulence term in pod models as exemplified for a high-reynolds-number flow over an ahmed body. *Journal of Fluid Mechanics*, 747:518–544, 2014.

F. Pedregosa and G. et al. Varoquaux. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12, 2011.

R Perez, G Bartram, T Beberniss, R Wiebe, and SM Spottswood. Calibration of aero-structural reduced order models using full-field experimental measurements. *Mechanical Systems and Signal Processing*, 86:49–65, 2017.

L. Perret, E. Collin, and J. Delville. Polynomial identification of POD based low-order dynamical system. *Journal of Turbulence*, (7):N17, 2006.

Robert N. C. Pfeifer, Jutho Haegeman, and Frank Verstraete. Faster identification of optimal contraction sequences for tensor networks. *Phys. Rev. E*, 90:033315, Sep 2014. . URL `https://link.aps.org/doi/10.1103/PhysRevE.90.033315`.

M. Poliashenko and C. K Aidun. A direct method for computation of simple bifurcations. *Journal of Computational Physics*, 121(2):246–260, 1995.

S. B. Pope. Turbulent flows, 2001.

C. Ramirez, V. Kreinovich, and M. Argaez. Why l_1 is a good approximation to l_0: A geometric explanation. *Journal of Uncertain Systems*, 7(3):203–207, 2013.

Sivaguru S Ravindran. A reduced-order approach for optimal control of fluids using proper orthogonal decomposition. *International journal for numerical methods in fluids*, 34(5):425–448, 2000.

D. Rempfer and H. F Fasel. Dynamics of three-dimensional coherent structures in a flat-plate boundary layer. *Journal of Fluid Mechanics*, 275:257–283, 1994a.

D. Rempfer and H. F. Fasel. Evolution of three-dimensional coherent structures in a flat-plate boundary layer. *Journal of Fluid Mechanics*, 260:351–375, 1994b.

Gemma Roig, Xavier Boix, and Fernando De la Torre. Optimal feature selection for subspace image matching. In *Computer Vision Workshops (ICCV Workshops), 2009 IEEE 12th International Conference on*, pages 200–205. IEEE, 2009.

Ribana Roscher, Bastian Bohn, Marco F Duarte, and Jochen Garcke. Explainable machine learning for scientific insights and discoveries. *Ieee Access*, 8:42200–42216, 2020.

C. W. Rowley and S. T. M. Dawson. Model reduction for flow analysis and control. *Annual Review of Fluid Mechanics*, 49:387–417, 2017.

C. W. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D. S. Henningson. Spectral analysis of nonlinear flows. *Journal of Fluid Mechanics*, 641:115–127, December 2009.

Riccardo Rubini, Davide Lasagna, and Andrea Da Ronch. l1-based calibration of pod-galerkin models of two-dimensional unsteady flows. *Chinese Journal of Aeronautics*, 2020a. ISSN 1000-9361. . URL `http://www.sciencedirect.com/science/article/pii/S1000936120305094`.

Riccardo Rubini, Davide Lasagna, and Andrea Da Ronch. l_1-based sparsification of energy interactions in two-dimensional turbulent flows. *arXiv preprint arXiv:2006.05736*, 2020b.

Riccardo Rubini, Davide Lasagna, and Andrea Da Ronch. A-priori sparsification of galerkin-based reduced order models. *arXiv preprint arXiv:2105.06753*, 2021.

M. Schlegel and B. R. Noack. On long-term boundedness of galerkin models. *Journal of Fluid Mechanics*, 765:325–352, 2015.

P. J. Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656:5–28, 2010.

M. Schmidt and H. Lipson. Distilling Free-Form Natural Laws from Experimental Data. *Science*, 324(5923):81–85, 2009.

Oliver T Schmidt. Bispectral mode decomposition of nonlinear flows. *arXiv preprint arXiv:2002.04146*, 2020.

M. Sieber, C. O. Paschereit, and K. Oberleithner. Spectral Proper Orthogonal decomposition. *Journal of Fluid Mechanics*, 792:798–828, 2016.

Witold Stankiewicz, Robert Roszak, and Marek Morzyński. Genetic algorithm-based calibration of reduced order galerkin models. *Mathematical Modelling and Analysis*, 16 (2):233–247, 2011. . URL `https://doi.org/10.3846/13926292.2011.579187`.

Krister Svanberg. A class of globally convergent optimization methods based on conservative convex separable approximations. *SIAM Journal on Optimization*, pages 555–573, 2014.

K. Taira, A. G. Nair, and S. L. Brunton. Network structure of two-dimensional decaying isotropic turbulence. *Journal of Fluid Mechanics*, 795, 2016.

K. Taira, S. L. Brunton, S. T. M. Dawson, C. W. Rowley, T. Colonius, B. J. McKeon, O. T. Schmidt, S. Gordeyev, V. Theofilis, and L. S Ukeiley. Modal analysis of fluid flows: An overview. *AIAA Journal*, pages 4013–4041, 2017.

Hendrik Tennekes and John Leask Lumley. *A first course in turbulence*. MIT press, 1972.

F. Terragni, E. Valero, and J. M. Vega. Local POD plus galerkin projection in the unsteady lid-driven cavity problem. *SIAM Journal on Scientific Computing*, 33(6):3538–3561, 2011.

V. L. Thomas, B. K. Lieu, M. R. Jovanović, B. F. Farrell, P. J. Ioannou, and D. F. Gayme. Self-sustaining turbulence in a restricted nonlinear model of plane Couette flow. *Physics of Fluids*, 26(10):105112, 2014.

V. L. Thomas, B. F. Farrell, P. J. Ioannou, and D. F. Gayme. A minimal model of self-sustaining turbulence. *Physics of Fluids*, 27(10):105104, 2015.

R. Tibshirani. Regression shrinkage and selection via the LASSO. *Journal of the Royal Statistical Society, B*, 58(1), 1996.

R. J. Tibshirani. The LASSO problem and uniqueness. *Electronic Journal of Statistics*, 7: 1456–1490, 2013.

Timo Tonn, Karsten Urban, and Stefan Volkwein. Optimal control of parameter-dependent convection-diffusion problems around rigid bodies. *SIAM Journal on scientific computing*, 32(3):1237–1260, 2010.

A. Towne, O. T. Schmidt, and T. Colonius. Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis. *Journal of Fluid Mechanics*, 847:821–867, 2018.

X. Xie, M. Mohebujjaman, L. G. Rebholz, and T. Iliescu. Data-Driven Filtered Reduced Order Modeling of Fluid Flows. *SIAM Journal of Scientific Computing*, 40(3):B834–B857, 2018.

L. Zhang and H. Schaeffer. On the convergence of the SINDy algorithm. *Multiscale Modeling & Simulation*, 17(3):948–972, 2019.

Wei Zhang and Mingjun Wei. Model order reduction using dmd modes and adjoint dmd modes. In *8th AIAA Theoretical Fluid Mechanics Conference*, page 3482, 2017.